

A KNOWLEDGE-INTENSIVE MACHINE-LEARNING APPROACH
TO THE PRINCIPAL-AGENT PROBLEM

By

KIRAN K. GARIMELLA

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA
1993

To my mother, Dr. Seeta Garimella

ACKNOWLEDGMENTS

I thank Prof. Gary Koehler, chairman of the DIS department, a guru to me in the deepest sense of the word who made it possible for me to grow intellectually and experience the richness and fulfillment of an active mind.

I also want to thank Prof. Selcuk Erenguc for encouraging me at all times; Prof. Harold Benson who taught me care, caution, and clarity in thinking by patiently teaching me proof techniques in mathematics; Prof. David E.M. Sappington for giving me invaluable lessons, by his teaching and example, on research techniques, for writing papers and books that are replete with elegance and clarity, and for ensuring that my research is meaningful and interesting from an economist's perspective; Prof. Sanford V. Berg, for providing valuable suggestions in agency theory; and Prof. Richard Elnicki, Prof. Antal Majthay, and Prof. Ira Horowitz for their advice and help with the research.

I thank Prof. Malay Ghosh, Department of Statistics, and Prof. Scott McCullough, Department of Mathematics, for their guidance in statistics and mathematics.

I also thank the administrative staff of the DIS department for helping me in numerous ways and making my work extremely pleasant.

I thank my wife, Raji, for her patience and understanding while I put in long and erratic hours.

I cannot conclude without expressing my deepest sense of gratitude to my mother, Dr. Seeta Garimella, who constantly encouraged me in ways too numerous to recount and made it possible for me to pursue my studies in the land of my dreams.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	iii
LIST OF TABLES	viii
ABSTRACT	xv
1 OVERVIEW	1
2 EXPERT SYSTEMS AND MACHINE LEARNING	6
2.1 Introduction	6
2.2 Expert Systems	8
2.3 Machine Learning	10
2.3.1 Introduction	10
2.3.2 Definitions and Paradigms	14
2.3.3 Probably Approximately Close Learning	21
3 GENETIC ALGORITHMS	23
3.1 Introduction	23
3.2 The Michigan Approach	26
3.3 The Pitt Approach	27
4 THE MAXIMUM ENTROPY PRINCIPLE	28
4.1 Historical Introduction	28
4.2 Examples	34
5 THE PRINCIPAL-AGENT PROBLEM	38
5.1 Introduction	38
5.1.1 The Agency Relationship	38
5.1.2 The Technology Component of Agency	40
5.1.3 The Information Component of Agency	40
5.1.4 The Timing Component of Agency	42

5.1.5	Limited Observability, Moral Hazard, and Monitoring	44
5.1.6	Informational Asymmetry, Adverse Selection, and Screening	45
5.1.7	Efficiency of Cooperation and Incentive Compatibility	47
5.1.8	Agency Costs	47
5.2	Formulation of the Principal-Agent Problem	48
5.3	Main Results in the Literature	62
5.3.1	Model 1: The Linear-Exponential-Normal Model	63
5.3.2	Model 2	68
5.3.3	Model 3	72
5.3.4	Model 4: Communication under Asymmetry	77
5.3.5	Model G: Some General Results	80
6	METHODOLOGICAL ANALYSIS	82
7	MOTIVATION THEORY	87
8	RESEARCH FRAMEWORK	92
9	MODEL 3	97
9.1	Introduction	97
9.2	An Implementation and Study	101
9.3	Details of Experiments	106
9.3.1	Rule Representation	106
9.3.2	Inference Method	110
9.3.3	Calculation of Satisfaction	111
9.3.4	Genetic Learning Details	114
9.3.5	Statistics Captured for Analysis	115
9.4	Results	116
9.5	Analysis of Results	118
10	REALISTIC AGENCY MODELS	149
10.1	Characteristics of Agents	157
10.2	Learning with Specialization and Generalization.	158
10.3	Notation and Conventions	160
10.4	Model 4: Discussion of Results	161
10.5	Model 5: Discussion of Results	163
10.6	Model 6: Discussion of Results	164
10.7	Model 7: Discussion of Results	165
10.8	Comparison of the Models	167
10.9	Examination of Learning	172
11	CONCLUSION	194

12	FUTURE RESEARCH	198
12.1	Nature of the Agency	198
12.2	Behavior and Motivation Theory	199
12.3	Machine Learning	200
12.4	Maximum Entropy	203
APPENDIX	FACTOR ANALYSIS	204
REFERENCES	206
BIOGRAPHICAL SKETCH	219

LIST OF TABLES

<u>Table</u>	<u>page</u>
9.1: Characterization of Agents	125
9.2: Iteration of First Occurrence of Maximum Fitness	126
9.3: Learning Statistics for Fitness of Final Knowledge Bases	126
9.4: Entropy of Final Knowledge Bases and Closeness to the Maximum	126
9.5: Frequency (as Percentage) of Values of Compensation Variables in the Final Knowledge Base in Experiment 1	127
9.6: Range, Mean and Standard Deviation of Values of Compensation Variables in the Final Knowledge Base in Experiment 1	127
9.7: Correlation Analysis of Values of Compensation Variables in the Final Knowledge Base in Experiment 1	128
9.8: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 1	128
9.9: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 1 - Factor Pattern	129
9.10: Experiment 1 - Varimax Rotation	130
9.11: Frequency (as Percentage) of Values of Compensation Variables in the Final Knowledge Base in Experiment 2	131
9.12: Range, Mean and Standard Deviation of Values of Compensation Variables in the Final Knowledge Base in Experiment 2	131
9.13: Correlation Analysis of Values of Compensation Variables in the Final Knowledge Base in Experiment 2	131

9.14: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 2 - Eigenvalues of the Correlation Matrix	132
9.15: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 2 - Factor Pattern	133
9.16: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 2 - Varimax Rotated Factor Pattern	134
9.17: Frequency (as Percentage) of Values of Compensation Variables in the Final Knowledge Base in Experiment 3	135
9.18: Range, Mean and Standard Deviation of Values of Compensation Variables in the Final Knowledge Base in Experiment 3	135
9.19: Correlation Analysis of Values of Compensation Variables in the Final Knowledge Base in Experiment 3	135
9.20: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 3 - Eigenvalues of the Correlation Matrix	136
9.21: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 3 - Factor Pattern	137
9.22: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 3 - Varimax Rotated Factor Pattern	138
9.23: Frequency (as Percentage) of Values of Compensation Variables in the Final Knowledge Base in Experiment 4	139
9.24: Range, Mean and Standard Deviation of Values of Compensation Variables in the Final Knowledge Base in Experiment 4	139
9.25: Correlation Analysis of Values of Compensation Variables in the Final Knowledge Base in Experiment 4	139
9.26: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 4 - Eigenvalues of the Correlation Matrix	140
9.27: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 4 - Factor Pattern	141
9.28: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 4 - Varimax Rotated Factor Pattern	143

9.29: Frequency (as Percentage) of Values of Compensation Variables in the Final Knowledge Base in Experiment 5	144
9.30: Range, Mean and Standard Deviation of Values of Compensation Variables in the Final Knowledge Base in Experiment 5	144
9.31: Correlation Analysis of Values of Compensation Variables in the Final Knowledge Base in Experiment 5	144
9.32: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 5 - Eigenvalues of the Correlation Matrix	145
9.33: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 5 - Factor Pattern	145
9.34: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 5 - Varimax Rotated Factor Pattern	146
9.35: Summary of Factor Analytic Results for the Five Experiments	146
9.36: Expected Factor Identification of Compensation Variables for the Five Experiments Derived from the Direct Factor Analytic Solution	147
9.37: Expected Factor Identification of Compensation Variables for the Five Experiments Derived from the Varimax Rotated Factor Analytic Solution	147
9.38: Expected Factor Identification of Behavioral and Risk Variables for the Five Experiments Derived from the Direct Factor Pattern	148
9.39: Expected Factor Identification of Behavioral and Risk Variables for the Five Experiments Derived from Varimax Rotated Factor Analytic Solution	148
10.1: Correlation of LP and CP with Simulation Statistics (Model 4)	174
10.2: Correlation of LP and CP with Compensation Offered to Agents (Model 4)	174
10.3: Correlation of LP and CP with Compensation in the Principal's Final KB (Model 4)	174
10.4: Correlation of LP and CP with the Movement of Agents (Model 44)	174

10.5: Correlation of LP with Agent Factors (Model 4)	174
10.6: Correlation of LP and CP with Agents' Satisfaction (Model 4)	175
10.7: Correlation of LP and CP with Agents' Satisfaction at Termination (Model 4)	175
10.8: Correlation of LP and CP with Agency Interactions (Model 4)	175
10.9: Correlation of LP with Rule Activation (Model 4)	175
10.10: Correlation of LP with Rule Activation in the Final Iteration (Model 4)	175
10.11: Correlation of LP and CP with Principal's Satisfaction and Least Squares (Model 4)	175
10.12: Correlation of Agent Factors with Agent Satisfaction (Model 4)	176
10.13: Correlation of Principal's Satisfaction with Agent Factors (Model 4)	176
10.14: Correlation of Principal's Satisfaction with Agents' Satisfaction (Model 4)	176
10.15: Correlation of Principal's Last Satisfaction with Agents' Last Satisfaction (Model 4)	176
10.16: Correlation of Principal's Factor with Agent Factors (Model 4)	177
10.17: Correlation of LP and CP with Simulation Statistics (Model 5)	177
10.18: Correlation of LP and CP with Compensation Offered to Agents (Model 5)	177
10.19: Correlation of LP and CP with Compensation in the Principal's Final Knowledge Base (Model 5)	177
10.20: Correlation of LP and CP with the Movement of Agents (Model 5)	177
10.21: Correlation of LP with Agent Factors (Model 5)	178
10.22: Correlation of LP and CP with Agents' Satisfaction (Model 5)	178
10.23: Correlation of LP and CP with Agents' Satisfaction at Termination (Model 5)	178

10.24: Correlation of LP and CP with Agency Interactions (Model 5)	178
10.25: Correlation of LP with Rule Activation (Model 5)	178
10.26: Correlation of LP with Rule Activation in the Final Iteration (Model 5 . .	179
10.27: Correlation of LP and CP with Payoffs from Agents (Model 5)	179
10.28: Correlation of LP and CP with Principal's Satisfaction, Principal's Factor and Least Squares (Model 5)	179
10.29: Correlation of Agent Factors with Agent Satisfaction (Model 5)	179
10.30: Correlation of Principal's Satisfaction with Agent Factors (Model 5) . . .	180
10.31: Correlation of Principal's Satisfaction with Agents' Satisfaction (Model 5)	180
10.32: Correlation of Principal's Last Satisfaction with Agents' Last Satisfaction (Model 5)	180
10.33: Correlation of Principal's Satisfaction with Outcomes from Agents (Model 5)	181
10.34: Correlation of Principal's Factor with Agents' Factors (Model 5)	181
10.35: Correlation of LP and CP with Simulation Statistics (Model 6)	181
10.36: Correlation of LP and CP with Compensation Offered to Agents (Model 6)	181
10.37: Correlation of LP and CP with Compensation in the Principal's Final Knowledge Base (Model 6)	182
10.38: Correlation of LP and CP with the Movement of Agents (Model 6)	182
10.39: Correlation of LP and CP with Agent Factors (Model 6)	182
10.40: Correlation of LP and CP with Agents' Satisfaction (Model 6)	182
10.41: Correlation of LP and CP with Agents' Satisfaction at Termination (Model 6)	183
10.42: Correlation of LP and CP with Agency Interactions (Model 6)	183

10.43: Correlation of LP and CP with Rule Activation (Model 6)	183
10.44: Correlation of LP and CP with Rule Activation in the Final Iteration (Model 6)	183
10.45: Correlation of LP and CP with Principal's Satisfaction and Least Squares (Model 6)	184
10.46: Correlation of Agents' Factors with Agents' Satisfaction (Model 6)	184
10.47: Correlation of Principal's Satisfaction with Agents' Factors and Agents' Satisfaction (Model 6)	185
10.48: Correlation of Principal's Factor with Agents' Factor (Model 6)	185
10.49: Correlation of LP and CP with Simulation Statistics (Model 7)	185
10.50: Correlation of LP and CP with Compensation Offered to Agents (Model 7)	185
10.51: Correlation of LP and CP with Compensation in the Principal's Final Knowledge Base (Model 7)	186
10.52: Correlation of LP and CP with the Movement of Agents (Model 7)	186
10.53: Correlation of LP with Agent Factors (Model 7)	186
10.54: Correlation of LP and CP with Agents' Satisfaction (Model 7)	186
10.55: Correlation of LP and CP with Agents' Satisfaction at Termination (Model 7)	187
10.56: Correlation of LP and CP with Agency Interactions (Model 7)	187
10.57: Correlation of LP and CP with Rule Activation (Model 7)	187
10.58: Correlation of LP with Rule Activation in the Final Iteration (Model 7) . .	187
10.59: Correlation of LP and CP with Payoffs from Agents (Model 7)	188
10.60: Correlation of LP and CP with Principal's Satisfaction (Model 7)	188
10.61: Correlation of Agent Factors with Agent Satisfaction (Model 7)	188

10.62: Correlation of Principal's Satisfaction with Agent Factors (Model 7) . . .	188
10.63: Correlation of Principal's Satisfaction with Agents' Satisfaction (Model 7)	189
10.64: Correlation of Principal's Last Satisfaction with Agents' Last Satisfaction (Model 7)	189
10.65: Correlation of Principal's Satisfaction with Outcomes from Agents (Model 7)	189
10.66: Correlation of Principal's Factor with Agents' Factor (Model 7)	189
10.67: Comparison of Models	190
10.68: Probability Distributions for Models 4, 5, 6, and 7	193

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

A KNOWLEDGE-INTENSIVE MACHINE-LEARNING APPROACH
TO THE PRINCIPAL-AGENT PROBLEM

By

Kiran K. Garimella

August 1993

Chairperson: Gary J. Koehler

Major Department: Decision and Information Sciences

The objective of the research is to explore an alternative approach to the solution of the principal-agent problem, which is extremely important since it is applicable in almost all business environments. It has been traditionally addressed by the optimization-analytical framework. However, there is a clearly recognized need for techniques that allow the incorporation of behavioral and motivational characteristics of the agent and the principal that influence their selection of effort and payment levels.

The alternative proposed is a knowledge-intensive, machine-learning approach, where all the relevant knowledge and the constraints of the problem are taken into account in the form of knowledge-bases.

Genetic algorithms are employed for learning, supplemented in later models by specialization and generalization operators. A number of models are studied in order of increasing complexity and realism. Initial studies are presented that provide counter-

examples to traditional agency theory and that emphasize the need for going beyond the traditional framework. The new framework is more robust, easily extensible in a modular manner, and yields contracts tailored to the behavioral characteristics of individual agents.

Factor analysis of final knowledge bases after extensive learning shows that elements of compensation besides basic pay and share of output play a greater role in characterizing good contracts. The learning algorithms tailor contracts to the behavioral and motivational characteristics of individual agents. Further, neither did perfect information yield the highest satisfaction nor did the complete absence of information yield the least satisfaction. This calls into question the traditional agency wisdom that more information is always desirable.

Studies of other models study the effect of two different policies of evaluating agents' performance by the principal-individualized (discriminatory) evaluation versus the relative (nondiscriminatory) evaluation. The results suggest guidelines for employing different types of models to simulate different agency environments.

CHAPTER 1 OVERVIEW

The basic research addressed by this dissertation is the theory and application of machine learning to assist in the solution of decision problems in business. Much of the earlier research in machine learning was devoted to addressing specific and ad-hoc problems or to fill a gap or make up for some deficiency in an existing framework, usually motivated by developments in expert systems and statistical pattern recognition. The first applications were to technical problems such as knowledge acquisition, coping with a changing environment and filtering of noise (where filtering and optimal control were considered inadequate because of poorly understood domains), data or knowledge reduction (where the usual statistical theory is inadequate to express the symbolic richness of the underlying domain), and scene and pattern analysis (where the classical statistical techniques fail to take into account pertinent prior information; see for example, Jaynes, 1986a).

The initial research was concerned with gaining an understanding of learning in extremely simple toy world models, such as checkers (Samuel, 1963), SHRDLU blocks world (Winograd, 1972), and various discovery systems. The insights gained by such research soon influenced serious applications.

The underlying domains of most of the early applications were relatively well structured, whether they were the stylized rules of checkers and chess or the digitized images of visual sensors. Our research focus is on importing these ideas into the area of business decisionmaking.

Genetic algorithms, a relatively new paradigm of machine learning, deals with adaptive processes modeled on ideas from natural genetics. Genetic algorithms use the ideas of parallelism, randomized search, fitness criteria for individuals, and the formation of new exploratory solutions using reproduction, survival and mutation. The concept is extremely elegant, powerful, and easy to work with from the viewpoint of the amount of knowledge necessary to start the search for solutions.

A related issue is maximum entropy. The Maximum Entropy Principle is an extension of Bayesian theory and is founded on two other principles: the Desideratum of Consistency and Maximal-Noncommitment. While Bayesian analysis begins by assuming a prior, the Maximum Entropy Principle seeks distributions that maximize the Shannon entropy and at the same time satisfy whatever constraints may apply. The justification for using Shannon entropy comes from the works of Bernoulli, Laplace, Jeffreys, and Cox on the one hand, and from the works of Maxwell, Boltzmann, Gibbs, and Shannon on the other; the principle has been extensively championed by Jaynes and is only just now penetrating into economic analysis.

Under the maximum entropy technique, the task of updating priors based on data is now subsumed under the general goal of maximizing entropy of distributions given any and all applicable constraints, where the data (or sufficient statistics on the data) play the

role of constraints. Maximum entropy is related to machine learning by the fact that the initial distributions (or assumptions) used in a learning framework, such as genetic algorithms, may be maximum entropy distributions. A topic of research interest is the development of machine learning algorithms or frameworks that are robust with respect to maximum entropy. In other words, deviation of initial distributions from maximum entropy distributions should not have any significant effect on the learning algorithms (in the sense of departure from good solutions).

The overall goal of the research is to present an integrated methodology involving machine learning with genetic algorithms in knowledge bases and to illustrate its use by application to an important problem in business. The principal-agent problem was chosen for the following reasons: it is widespread, important, nontrivial, and fairly general so that different models of the problem can be investigated, and information-theoretic considerations play a crucial role in the problem. Moreover, a fair amount of interest over the problem has been generated among researchers in economics, finance, accounting, and game theory, whose predominant approach to the problem is that of constrained optimization. Several analytical insights have been generated, which should serve as points of comparison to results that are expected from our new methodology.

The most important component of the new proposed methodology is information in the form of knowledge bases, coupled with strength of performance of the individual pieces of knowledge. These knowledge bases, the associated strengths, their relation to one another, and their role in the scheme of things are derived from the individuals' prior knowledge and from the theory of human behavior and motivation. These knowledge

bases contain, for example, information about the agent's characteristics and pattern of behavior under different compensation schemes; in other words, they deal with the issues of hidden characteristics and induced effort or behavior. Given the expected behavior pattern of an agent, a related research issue is the study of the effect of using distributions that have maximum entropy with respect to the expected behavior.

Trial compensation schemes, which come from the specified knowledge bases, are presented to the agent(s). Upon acceptance of the contract and realization of the output, the actual performance of the agent (in terms of output or the total welfare) is evaluated, and the associated compensation schemes are assigned proportional credit. Periodically, iterations of the genetic algorithm will be used to create a new knowledge base that enriches the current one.

Chapter 2 begins with an introduction to artificial intelligence, expert systems, and machine learning. Chapter 3 describes genetic algorithms. Chapter 4 covers the origin of the Maximum Entropy Principle and its formulation. Chapter 5 deals with a survey of the principal-agent problem, where a few basic models are presented, along with some of the main results of the research.

Chapter 6 examines the traditional methodology used in attacking the principal-agent problem, and measures to cover the inadequacies are proposed. One of the basic assumptions of the economic theory--the assumption of risk attitudes and utility--is circumvented by directly dealing with the knowledge-based models of the agent and the principal. To this end, a brief look at some of the ideas from behavior and motivation theory is taken in Chapter 7.

Chapter 8 describes the basic research model. Elements of behavior and motivation theory and knowledge bases are incorporated. A research strategy to study agency problems is proposed. The use of genetic algorithms periodically to enrich the knowledge bases and to carry out learning is suggested. An overview of the research models, all of which incorporate many features of the basic model, is presented.

Chapter 9 describes Model 3 in detail. Chapter 10 introduces Models 4 through 7 and describes each in detail. Chapter 11 provides a summary of the results of Chapters 9 and 10. Directions for future research are covered in Chapter 12.

CHAPTER 2 EXPERT SYSTEMS AND MACHINE LEARNING

2.1 Introduction

The use of artificial intelligence in a computerized world is as revolutionary as the use of computers is in a manual world. One can make computers intelligent in the same sense as man is intelligent. The various techniques of doing this compose the body of the subject of artificial intelligence. At the present state of the art, computers are at last being designed to compete with man on his own ground on something like equal terms. To put it in another way, computers have traditionally acted as convenient tools in areas where man is known to be deficient or inefficient, namely, doing complicated arithmetic very quickly, or making many copies of data (i.e., files, reports, etc.).

Learning new things, discovering facts, conjecturing, evaluating and judging complex issues (for example, consulting), using natural languages, analyzing and understanding complex sensory inputs such as sound and light, and planning for future action are mental processes that are peculiar to man (and to a lesser extent, to some animals). Artificial intelligence is the science of simulating or mimicking these mental processes in a computer.

The benefits are immediately obvious. First, computers already fill some of the gaps in human skills; second, artificial intelligence fills some of the gaps that computers

themselves suffer (i.e., human mental processes). While the full simulation of the human brain is a distant dream, limited application of this idea has already produced favorable results.

Speech-understanding problems were investigated with the help of the HEARSAY system (Erman et al., 1980, 1981; and Hayes-Roth and Lesser, 1977). The faculty of vision relates to pattern recognition and classification and analysis of scenes. These problems are especially encountered in robotics (Paul, 1981). Speech recognition coupled with natural language understanding as in the limited system SHRDLU (Winograd, 1973) can find immediate uses in intelligent secretary systems that can help in data management and correspondence associated with business.

An area that is commercially viable in large business environments that involve manufacturing and any other physical treatment of objects is robotics. This is a proven area of artificial intelligence application, but is not yet cost effective for small business. Several robot manufacturers have a good order book position. For a detailed survey see for example, Engelberger, 1980.

An interesting viewpoint to the application of artificial intelligence to industry and business is that presented by decision analysis theory. Decision analysis helps managers to decide between alternative options and assess risk and uncertainty in a better way than before, and to carry out conflict management when there are conflicts among objectives. Certain operations research techniques are also incorporated, as for example, fair allocation of resources that optimize returns. Decision analysis is treated in Fishburn (1981), Lindley (1971), Keeney (1984) and Keeney and Raiffa (1976). In most

applications of expert systems, concepts of decision analysis find expression (Phillips, 1986). Manual application of these techniques is not cost effective, whereas their use in certain expert systems, which go by the generic name of Decision Analysis Expert Systems, leads to quick solutions of what were previously thought to be intractable problems (Conway, 1986). Several systems have been proposed that range from scheduling to strategy planning. See for example, Williams (1986).

2.2 Expert Systems

The most fascinating and economically justifiable area of artificial intelligence is the development of expert systems. These are computer systems that are designed to provide expert advice in any area. The kind of information that distinguishes an expert from a nonexpert forms the central idea in any expert system. This is perhaps the only area that provides concrete and conclusive proof of the power of artificial intelligence techniques. Many expert systems are commercially viable and motivate diverse sources of funding for research into artificial intelligence. An expert system incorporates many of the techniques of artificial intelligence, and a positive response to artificial intelligence depends on the reception of expert systems by informed laymen.

To construct an expert system, the knowledge engineer works with an expert in the domain and extracts knowledge of relevant facts, rules, rules-of-thumb, exceptions to standard theory, and so on. This is a difficult task and is known variously as knowledge acquisition or mining. Because of the complex nature of the knowledge and the ways humans store knowledge, this is bound to be a bottleneck to the development

of the expert system. This knowledge is codified in the form of several rules and heuristics. Validation and verification runs are conducted on problems of sufficient complexity to see that the expert system does indeed model the thinking of the expert. In the task of building expert systems, the knowledge engineer is helped by several tools, such as EMYCIN, EXPERT, OPS5, ROSIE, GURU, etc.

The net result of the activity of knowledge mining is a knowledge base. An inference system or engine acts on this knowledge base to solve problems in the domain of the expert system. An important characteristic of expert systems is the ability to justify and explain their line of reasoning. This is to create credibility during their use. In order to do this, they must have a reasonably sophisticated input/output system.

Some of the typical problems handled by expert systems in the areas of business, industry, and technology are presented in Feigenbaum and McCorduck (1983) and Mitra (1986). Important cases where expert systems are brought in to handle the problems are

1. Capturing, replicating, and distributing expertise.
2. Fusing the knowledge of many experts.
3. Managing complex problems and amplifying expertise.
4. Managing knowledge.
5. Gaining a competitive edge.

As examples of successful expert systems, one can consider MYCIN, designed to diagnose infectious diseases (Shortliffe, 1976); DENDRAL, for interpretation of molecular spectra (Buchanan and Feigenbaum, 1978); PROSPECTOR, for geological studies (Duda et al., 1979; Hart, 1978); and WHY, for teaching geography (Stevens and

Collins, 1977). For a more exhaustive treatment, see, for example Stefik et al. (1982), Barr and Feigenbaum (1981, 1982), Cohen and Feigenbaum (1982), and Barr et al. (1989).

2.3 Machine Learning

2.3.1 Introduction

One of the key limitations of computers as envisaged by early researchers is the fact that they must be told in explicit detail how to solve every problem. In other words, they lack the capacity to learn from experience and improve their performance with time. Even in most expert systems today, there is only some weak form of implicit learning, such as learning by being told, rote memorizing, and checking for logical consistency. The task of machine learning research is to make up for this inadequacy by incorporating learning techniques into computers.

The abstract goals of machine learning research are broadly

1. To construct learning algorithms that enable computers to learn.
2. To construct learning algorithms that enable computers to learn in the same way as humans learn.

In both cases, the functional goals of machine learning research are as follows:

1. To use the learning algorithms in application domains to solve nontrivial problems.
2. To gain a better understanding of how humans learn, and the details of human cognitive processes.

When the goal is to come up with paradigms that can be used to solve problems, several subsidiary goals can be proposed:

1. To see if the learning algorithms do indeed perform better than humans do in similar situations.
2. To see if the learning algorithms come up with solutions that are intuitively meaningful for humans.
3. To see if the learning algorithms come up with solutions that are in some way better or less expensive than some alternative methodology.

It is undeniable that humans possess cognitive skills that are superior not only to other animals but also to most learning algorithms that are in existence today. It is true that some of these algorithms perform better than humans in some limited and highly formalized situations involving carefully modeled problems, just as the simplex method consistently produces solutions superior to those possible by a human being. However, and this is the crucial issue, humans are quick to adopt different strategies and solve problems that are ill-structured, ill-defined, and not well understood, for which there does not exist any extensive domain theory, and that are characterized by uncertainty, noise, or randomness. Moreover, in many cases, it seems more important to humans to find solutions to problems that satisfy some constraints rather than to optimize some "function." At the present state of the art, we do not have a consistent, coherent and systematic theory of what these constraints are. These constraints are usually understood to be behavioral or motivational in nature.

Recent research has shown that it is also undeniable that humans perform very poorly in the following respects:

- * they do not solve problems in probability theory correctly ;
- * while they are good at deciding cogency of information, they are poor at judging relevance (see Raiffa, accident witnesses, etc.);
- * they lack statistical sophistication;
- * they find it difficult to detect contradictions in long chains of reasoning;
- * they find it difficult to avoid bias in inference and in fact may not be able to identify it.

(See for example, Einhorn, 1982; Kahneman and Tversky, 1982a, 1982b, 1982c, 1982d; Lichtenstein et al., 1982; Nisbett et al., 1982; Tversky and Kahneman, 1982a, 1982b, 1982c, 1982d.)

Tversky and Kahneman (1982a) classify, for example, several misconceptions in probability theory as follows:

- * insensitivity to prior probability of outcomes;
- * insensitivity to sample size;
- * misconceptions of chance;
- * insensitivity to predictability;
- * the illusion of validity;
- * misconceptions of regression.

The above inadequacies on the part of humans pertain to higher cognitive thinking. It goes without saying that humans are poor at manipulating numbers quickly, and are subject to physical fatigue and lack of concentration when involved in mental activity for a long time. Computers are, of course, subject to no such limitations.

It is important to note that these inadequacies usually do not lead to disastrous consequences in most everyday circumstances. However, the complexity of the modern world gives rise to intricate and substantial problems, solutions to which forbid inadequacies of the above type.

Machine learning must be viewed as an integrated research area that seeks to understand the learning strategies employed by humans, incorporate them into learning algorithms, remove any cognitive inadequacies faced by humans, investigate the possibility of better learning strategies, and characterize the solutions yielded by such research in terms of proof of correctness, convergence to optimality (where meaningful), robustness, graceful degradation, intelligibility, credibility, and plausibility.

Such an integrated view does not see the different goals of machine learning research as separate and clashing; insights in one area have implications for another. For example, insights into how humans learn help spot their strengths and weaknesses, which motivates research into how to incorporate the strengths into algorithms and how to cover up the weaknesses; similarly, discovering solutions from machine learning algorithms that are at first nonintuitive to humans motivates deeper analysis of the domain theory and of the human cognitive processes in order to come up with at least plausible explanations.

2.3.2 Definitions and Paradigms

Any activity that improves performance or skills with time may be defined as learning. This includes motor skills and general problem-solving skills. This is a highly functional definition of learning and may be objected to on the grounds that humans learn even in a context that does not demand action or performance. However, the functional definition may be justified by noting that performance can be understood as improvement in knowledge and acquisition of new knowledge or cognitive skills that are potentially usable in some context to improve actions or enable better decisions to be taken.

Learning may be characterized by several criteria. Most paradigms fall under more than one category. Some of these are

1. Involvement of the learner.
2. Sources of knowledge.
3. Presence and role of a teacher.
4. Access to an oracle (learning from internally generated examples).
5. Learning "richness."
6. Activation of learning:
 - (a) systematic;
 - (b) continuous;
 - (c) periodic or random;
 - (d) background;
 - (e) explicit or external (also known as intentional);

- (f) implicit (also known as incidental);
- (g) call on success; and
- (h) call on failure.

When classified by the criterion of the learner's involvement, the standard is the degree of activity or passivity of the learner. The following paradigms of learning are classified by this criterion, in increasing order of learner control:

1. Learning by being told (learner only needs to memorize by rote);
2. Learning by instruction (learner needs to abstract, induce, or integrate to some extent, and then store it);
3. Learning by examples (learner needs to induce to a great extent the correct concept, examples of which are supplied by the instructor);
4. Learning by analogy (learner needs to abstract and induce to a greater degree in order to learn or solve a problem by drawing the analogy. This implies that the learner already has a store of cases against which he can compare the analogy and that he knows how to abstract and induce knowledge);
5. Learning by observation and discovery (here the role of the learner is greatest; the learner needs to focus on only the relevant observations, use principles of logic and evidence, apply some value judgments, and discover new knowledge either by using induction or deduction).

The above learning paradigms may also be classified on the basis of richness of knowledge. Under this criterion, the focus is on the richness of the resulting knowledge, which may be independent of the involvement of the learner. The spectrum of learning

is from "raw data" to simple functions, complicated functions, simple rules, complex knowledge bases, semantic nets, scripts, and so on.

One fundamental distinction can be made from observation of human learning. The most widespread form of human learning is incidental learning. The learning process is incidental to some other cognitive process. Perception of the world, for example, leads to formation of concepts, classification of objects in classes or primitives, the discovery of the abstract concepts of number, similarity, and so on (see for example, Rand 1967). These activities are not indulged in deliberately. As opposed to incidental learning, we have intentional learning, where there is a deliberate and explicit effort to learn. The study of human learning processes from the standpoint of implicit or explicit cognition is the main subject of research in psychological learning. (See for example, Anderson, 1980; Craik and Tulving, 1975; Glass and Holyoak, 1986; Hasher and Zacks, 1979; Hebb, 1961; Mandler, 1967; Reber, 1967; Reber, 1976; Reber and Allen, 1978; Reber et al., 1980).

A useful paradigm for the area of expert systems might be learning through failure. The explanation facility ensures that the expert system knows why it is correct when it is correct, but it needs to know why it is wrong when it is wrong, if it must improve performance with time. Failure analysis helps in focussing on deficient areas of knowledge.

Research in machine learning raises several wider epistemological issues such as hierarchy of knowledge, contextuality, integration, conditionality, abstraction, and reduction. The issue of hierarchy arises in induction of decision trees (see for example,

Quinlan, 1979; Quinlan, 1986; Quinlan, 1990); contextuality arises in learning semantics, as in conceptual dependency (see for example, Schank, 1972; Schank and Colby, 1973), learning by analogy (see for example, Buchanan et al., 1977; Dietterich and Michalski, 1979), and case-based reasoning (Riesbeck and Schank, 1989); integration is fundamental to forming relationships, as in semantic nets (Quillian, 1968; Anderson and Bower, 1973; Anderson, 1976; Norman, et al., 1975; Schank and Abelson, 1977), and frame-based learning (see for example, Minsky, 1975); abstraction deals with formation of universals or classes, as in classification (see for example, Holland, 1975), and induction of concepts (see for example, Mitchell, 1977; Mitchell, 1979; Valiant, 1984; Haussler, 1988); reduction arises in the context of deductive learning (see for example, Newell and Simon, 1956; Lenat, 1977), conflict resolution (see for example, McDermott and Forgy, 1978), and theorem-proving (see for example, Nilsson, 1980). For an excellent treatment of these issues from a purely epistemological viewpoint, see for example Rand (1967) and Peikoff (1991).

In discussing real-world examples of learning, it is difficult or meaningless to look for one single paradigm or knowledge representation scheme as far as learning is concerned. Similarly, there could be multiple teachers: humans, oracles, and an accumulated knowledge that acts as an internal generator of examples.

In analyzing learning paradigms, it is useful to look at least three aspects, since they each have a role in making the others possible:

1. Knowledge representation scheme.
2. Knowledge acquisition scheme.

3. Learning scheme.

At the present time, we do not yet have a comprehensive classification of learning paradigms and their systematic integration into a theory. One of the first attempts in this direction was taken by Michalski, Carbonell, and Mitchell (1983).

An extremely interesting area of research in machine learning that will have far-reaching consequences for such a theory of learning is multistrategy systems, which try to combine one or more paradigms or types of learning based on domain problem characteristics or to try a different paradigm when one fails. See for example Kodratoff and Michalski (1990). One may call this type of research meta-learning research, because the focus is not simply on rules and heuristics for learning, but on rules and heuristics for learning paradigms. Here are some simple learning heuristics, for example:

LH1: Given several "isa" relationships, find out about relations between the properties.

(For example, the observation that "Socrates is a man" motivates us to find out why Socrates should indeed be classified as a man, i.e., to discover that the common properties are "rational animal" and several physical properties.)

LH2: When an instance causes an existing heuristic with certainty to be revised downwards, ask for causes.

LH3: When an instance that was thought to belong to a concept or class but later turns out not to belong to it, find out what it does belong to.

LH4: If X isa Y1 and X isa Y2, then find the relationship between Y1 and Y2, and check for consistency. (This arises in learning by using semantic nets).

- LH5: Given an implication, find out if it is also an equivalence.
- LH6: Find out if any two or more properties are semantically the same, the opposite, or unrelated.
- LH7: If an object possesses two or more properties simultaneously from the same class or similar classes, check for contradictions, or rearrange classes hierarchically.
- LH8: An isa-tree in a semantic net creates an isa-tree with the object as a parent; find out in which isa-tree the parent object occurs as a child.

We can contrast these with meta-rules or meta-heuristics. A meta-rule is also a rule which says something about another rule. It is understood that meta-rules are watchdog rules that supervise the firing of other rules. Each learning paradigm has a set of rules that will lead to learning under that paradigm. We can have a set of meta-rules for learning if we have a learning system that has access to several paradigms of learning and if we are concerned with what paradigm to select at any given time. Learning meta-rules help the learner to pick a particular paradigm because the learner has knowledge of the applicability of particular paradigms given the nature and state of a domain or given the underlying knowledge-base representation schema.

The following are examples of meta-rules in learning:

- ML1: **If** several instances of a domain-event occur,
 then use generalization techniques.
- ML2: **If** an event or class of events occur a number of times with little or no change on
 each occurrence,
 then use induction techniques.

ML3: If a problem description similar to the problem on hand exists in a different domain or situation and that problem has a known solution,
 then use learning-by-analogy techniques.

ML4: If several facts are known about a domain including axioms and production rules,
 then use deductive learning techniques.

ML5: If undefined variables or unknown variables are present and no other learning rule was successful,
 then use the learning-from-instruction paradigm.

In all cases of learning, meta-rules dictate learning strategies, whether explicitly as in a multi-strategy system, or implicitly as when the researcher or user selects a paradigm.

Just as in expert systems, the learning strategy may be either goal directed or knowledge directed. Goal-directed learning proceeds as follows:

1. Meta-rules select learning paradigm(s).
2. Learner imposes the learning paradigm on the knowledge base.
3. The structure of the knowledge base and the characteristics of the paradigm determine the representation scheme.
4. The learning algorithm(s) of the paradigm(s) execute(s).

Knowledge directed learning, on the other hand, proceeds as follows:

1. The learner examines the available knowledge base.
2. The structure of the knowledge base limits the extent and type of learning, which is determined by the meta-rules.
3. The learner chooses an appropriate representation scheme.

4. The learning algorithm(s) of the chosen learning paradigm(s) execute(s).

2.3.3 Probably Approximately Close Learning

Early research on inductive inference dealt with supervised learning from examples (see for example, Michalski, 1983; Michalski, Carbonell, and Mitchell, 1983). The goal was to learn the correct concept by looking at both positive and negative examples of the concept in question. These examples were provided in one of two ways: either the learner obtained them by observation, or they were provided to the learner by some external instructor. In both cases, the class to which each example belonged was conveyed to the learner by the instructor (supervisor, or oracle). The examples provided to the learner were drawn from a population of examples or instances. This is the framework underlying early research in inductive inference (see for example, Quinlan, 1979; Quinlan, 1986; Angluin and Smith 1983).

Probably Approximately Close Identification (or PAC-ID for short) is a powerful machine-learning methodology that seeks inductive solutions in a supervised nonincremental learning environment. It may be viewed as a multiple-criteria learning problem in which there are at least three major objectives:

- (1) to derive (or induce) the correct solution, concept or rule, which is as close as we please to the optimal (which is unknown);
- (2) to achieve as high a degree of confidence as we please that the solution so derived above is in fact as close to the optimal as we intended;
- (3) to ensure that the "cost" of achieving the above two objectives is "reasonable."

PAC-ID therefore replaces the original research direction in inductive machine learning (seeking the true solution) by the more practical goal of seeking solutions close to the true one in polynomial time. The technique has been applied to certain classes of concepts, such as conjunctive normal forms (CNF). Estimates of necessary distribution independent sample sizes are derived based on the error and confidence criteria; the sample sizes are found to be polynomial in some factor such as the number of attributes. Applications to science and engineering have been demonstrated.

The pioneering work on PAC-ID was by Valiant (1984, 1985) who proposed the idea of finding approximate solutions in polynomial time. The ideas of characterizing the notion of approximation by using the concept of functional complexity of the underlying hypothesis spaces, introducing confidence in the closeness to optimality, and obtaining results that are independent of the underlying probability distribution with which the supervisory examples are generated (by nature or by the supervisor), compose the direction of the latest research. (See for example, Haussler, 1988; Haussler, 1990a; Haussler, 1990b; Angluin, 1987; Angluin, 1988; Angluin and Laird, 1988; Blumer, Ehrenfeucht, Haussler, and Warmuth, 1989; Pitt and Valiant, 1988; and Rivest, 1987).

The theoretical foundations for the mathematical ideas of learning convergence with high confidence are mainly derived from ideas in statistics, probability, statistical decision theory, and fractal theory. (See for example, Vapnik, 1982; Vapnik and Chervonenkis, 1971; Dudley, 1978; Dudley, 1984; Dudley, 1987; Kolmogorov and Tihomirov, 1961; Kullback, 1959; Mandelbrot, 1982; Pollard, 1984; Weiss and Kulikowski, 1991).

CHAPTER 3 GENETIC ALGORITHMS

3.1 Introduction

Genetic classification algorithms are learning algorithms that are modeled on the lines of natural genetics (Holland, 1975). Specifically, they use operators such as reproduction, crossover, mutation, and fitness functions. Genetic algorithms make use of inherent parallelism of chromosome populations and search for better solutions through randomized exchange of chromosome material and mutation. The goal is to improve the gene pool with respect to the fitness criterion from generation to generation.

In order to use the idea of genetic algorithms, problems must be appropriately modeled. The parameters or attributes that constitute an individual of the population must be specified. These parameters are then coded. The simulation begins with a random generation of an initial population of chromosomes, and the fitness of each is calculated. Depending on the problem and the type of convergence desired, it may be decided to keep the population size constant or varying across iterations of the simulation.

Using the population of an iteration, individuals are selected randomly according to their fitness level to survive intact or to mate with other similarly selected individuals. For mating members, a crossover point is randomly determined (an individual with n

attributes has $n-1$ crossover points), and the individuals exchange their "strings," thus forming new individuals. It may so happen that the new individuals are exactly the same as the parents. In order to introduce a certain amount of richness into the population, a mutation operator with extremely low probability is applied to the bits in the individual strings, which randomly changes each bit. After mating, survival, and mutation, the fitness of each individual in the new population is calculated. Since the probability of survival and mating is dependent on the fitness level, more fit individuals have a higher probability of passing on their genetic material.

Another factor plays a role in determining the average fitness of the population. Portions of the chromosome, called genes or features, act as determinants of qualities of the individual. Since in mating, the crossover point is chosen randomly, those genes that are shorter in length are more likely to survive a crossover and thus be carried from generation to generation. This has important implications for modeling a problem and will be mentioned in the chapter on research directions.

The power of genetic algorithms (henceforth, GAs) derives from the following features:

1. It is only necessary to know enough about the problem to identify the essential attributes of the solution (or "individual"); the researcher can work in comparative ignorance of the actual combinations of attribute values that may denote qualities of the individual.
2. Excessive knowledge cannot harm the algorithm; the simulation may be started with any extra knowledge the researcher may have about the problem,

such as his beliefs about which combinations play an important role. In such cases, the simulation may start with the researcher's population and not a random population; if it turns out that the whole or some part of this knowledge is incorrect or irrelevant, then the corresponding individuals get low fitness values and hence have a high probability of eventually disappearing from the population.

3. The remarks in point 2 above apply in the case of mutation also. If mutation gives rise to a useless feature, that individual gets a low fitness value and hence has a low probability of remaining in the population for a long time.
4. Since GAs use many individuals, the probability of getting stuck at local optima is minimized.

According to Holland (1975), there are essentially four ways in which genetic algorithms differ from optimization techniques:

1. GAs manipulate codings of attributes directly.
2. They conduct search from a population and not from a single point.
3. It is not necessary to know or assume extra simplifications in order to conduct the search; GAs conduct the search "blindly." It must be noted however, that randomized search does not imply directionless search.
4. The search is conducted using stochastic operators (random selection according to fitness) and not by using deterministic rules.

There are two important models for GAs in learning. One is the Pitt approach, and the other is the Michigan approach. The approaches differ in the way they define individuals and the goals of the search process.

3.2 The Michigan Approach

The knowledge base of the researcher or the user constitutes the genetic population, in which each rule is an individual. The antecedents and consequents of each rule form the chromosome. Each rule denotes a classifier or detector of a particular signal from the environment. Upon receipt of a signal, one or more rules fire, depending on the signal satisfying the antecedent clauses. Depending on the success of the action taken or the consequent value realized, those rules that contributed to the success are rewarded, and those rules that supported a different consequent value or action are punished. This process of assigning reward or punishment is called credit assignment.

Eventually, rules that are correct classifiers get high reward values, and their proposed action when fired carries more weight in the overall decision of selecting an action. The credit assignment problem is the problem of how to allocate credit (reward or punishment). One approach is the bucket-brigade algorithm (Holland, 1986).

The Michigan approach may be combined with the usual genetic operators to investigate other rules that may not have been considered by the researcher.

3.3 The Pitt Approach

The Pitt Approach, by De Jong (see for example, De Jong, 1988), considers the whole knowledge base as one individual. The simulation starts with a collection of knowledge bases. The operation of crossover works by randomly dichotomizing two parent knowledge bases (selected at random) and mixing the dichotomized portions across the parents to obtain two new knowledge bases. The Pitt approach may be used when the researcher has available to him a panel of experts or professionals, each of whom provides one knowledge base for some decision problem at hand. The crossover operator therefore enables one to consider combinations of the knowledge of the individuals, a process that resembles a brainstorming session. This is similar to a group decision-making approach. The final knowledge base or bases that perform well empirically would then constitute a collection of rules obtained from the best rules of the original expertise, along with some additional rules that the expert panel did not consider before. The Michigan approach will be used in this research to simulate learning on one knowledge base.

CHAPTER 4 THE MAXIMUM ENTROPY PRINCIPLE

4.1 Historical Introduction

The principle of maximum entropy was championed by E.T. Jaynes in the 1950s and has gained many adherents since. There are a number of excellent papers by E.T. Jaynes explaining the rationale and philosophy of the maximum entropy principle. The discussion of the principle essentially follows Jaynes (1982, 1983, 1986a, 1986b, and 1991).

The maximum entropy principle may be viewed as "a natural extension and unification of two separate lines of development. . . . The first line is identified with the names Bernoulli, Laplace, Jeffreys, Cox; the second with Maxwell, Boltzmann, Gibbs, Shannon." (Jaynes, 1983).

The question of approaching any decision problem with some form of prior information is historically known as the Principle of Insufficient Reason (so named by James Bernoulli in 1713). Jaynes (1983) suggests the name Desideratum of Consistency, which may be formally stated as follows:

- (1) a probability assignment is a way of describing a certain state of knowledge;
i.e., probability is an epistemological concept, not a metaphysical one;

- (2) when the available evidence does not favor any one alternative among others, then the state of knowledge is described correctly by assigning equal probabilities to all the alternatives;
- (3) suppose A is an event or occurrence for which some favorable cases out of some set of possible cases exist. Suppose also that all the cases are equally likely. Then, the probability that A will occur is the ratio of the number of cases favorable to A to the total number of equally possible cases. This idea is formally expressed as

$$Pr[A] = \frac{M}{N} = \frac{\text{Number of cases favorable to A}}{\text{Number of equally possible cases}}.$$

In cases where $Pr[]$ is difficult to estimate (such as when the number of cases is infinite or impossible to find out), Bernoulli's weak law of large numbers may be applied, where

$$\begin{aligned} Pr[A] &= \frac{M}{N} = \frac{\text{Number of cases favorable to A}}{\text{Total number of equally likely cases}} \\ &= \frac{\text{Number of times A occurs}}{\text{Number of trials}} \\ &= \frac{m}{n}. \end{aligned}$$

Limit theorems in statistics show that given (M, N) as the true state of nature, the observed frequency $f(m, n) = m/n$ approaches $Pr[A] = P(M, N) = M/N$ as the number of trials increase.

The reverse problem consists of estimating $P(M,N)$ by $f(m,n)$. For example, the probability of seeing m successes in n trials when each trial is independent with probability of success p , is given by the binomial distribution:

$$P(m \mid n, \frac{M}{N}) = P(m \mid n, p) = \binom{n}{m} p^m (1-p)^{n-m}.$$

The inverse problem would then consist of finding $\Pr[M]$ given (m,N,n) . This problem was given a solution by Bayes in 1763 as follows: Given (m,n) , then

$$\begin{aligned} \Pr[p < \frac{M}{N} < p + dp] &= P(dp \mid m, n) \\ &= \frac{(n+1)!}{m! (n-m)!} p^m (1-p)^{n-m} dp. \end{aligned}$$

which is the Beta distribution.

These ideas were generalized and put into the form they are today, known as the Bayes' theorem, by Laplace as follows: When there is an event E with possible causes C_i , and given prior information I and the observation E , the probability that a particular cause C_i caused the event E is given by

$$P(C_i \mid E, I) = \frac{P(E \mid C_i) P(C_i \mid I)}{\sum_j P(E \mid C_j) P(C_j \mid I)}$$

which result has been called "learning by experience" (Jaynes, 1978).

The contributions of Laplace were rediscovered by Jeffreys around 1939 and in 1946 by Cox who, for the first time, set out to study the "possibility of constructing a consistent set of mathematical rules for carrying out plausible, rather than deductive, reasoning." (Jaynes, 1983).

According to Cox, the fundamental result of mathematical inference may be described as follows: Suppose A, B, and C represent propositions, AB the proposition "Both A and B are true", and $\neg A$ the negation of A. Then, the consistent rules of combination are:

$$P(AB|C) = P(A|BC) P(B|C), \text{ and}$$

$$P(A|B) + P(\neg A|B) = 1.$$

Thus, "Cox proved that any method of inference in which we represent degrees of plausibility by real numbers, is necessarily either equivalent to Laplace's, or inconsistent." (Jaynes, 1983).

The second line of development starts with James Clerk Maxwell in the 1850s who, in trying to find the probability distribution for the velocity direction of spherical molecules after impact, realized that knowledge of the meaning of the physical parameters of any system constituted extremely relevant prior information. The development of the concept of entropy maximization started with Boltzmann who investigated the distribution of molecules in a conservative force field in a closed system.

Given that there are N molecules in the closed system, the total energy E remains constant irrespective of the distribution of the molecules inside the system. All positions and velocities are not equally likely. The problem is to find the most probable distribution of the molecules. Boltzmann partitioned the phase space of position and momentum into a discrete number of cells R_k , where $1 \leq k \leq s$. These cells were assumed to be such that the k-th cell is a region which is small enough so that the energy of a molecule as it moves inside that region does not change significantly, but which is

also so large that a large number N_k of molecules can be accommodated in it. The problem of Boltzmann then reduces to the problem of finding the best prediction of N_k for any given k in $1, \dots, s$.

The numbers N_k are called the occupation numbers. The number of ways a given set of occupation numbers will be realized is given by the multinomial coefficient

$$W(N_k) = \frac{N!}{N_1! N_2! \dots N_s!} \quad \dots (1)$$

The constraints are given by

$$E = \sum_{k=1}^s N_k E_k, \quad \text{and}$$

$$N = \sum_{k=1}^s N_k.$$

Since each set $\{N_k\}$ of occupation numbers represents a possible distribution, the problem is equivalently expressed as finding the most probable set of occupation numbers from the many possible sets. Using Stirling's approximation of factorials

$$n! \approx \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$$

in equation (1) yields

$$\log W = -N \sum_{k=1}^s \left(\frac{N_k}{N}\right) \log\left(\frac{N_k}{N}\right). \quad \dots (2)$$

The right hand side of (2) is the familiar Shannon entropy formula for the distribution specified by probabilities which are approximated by the frequencies N_k/N , $k = 1, \dots, s$. In fact, in the limit as N goes to infinity,

$$\lim_{N \rightarrow \infty} N^{-1} \log W = -\sum \frac{N_i}{N} \log \left(\frac{N_i}{N} \right) = H.$$

Distributions of higher entropy therefore have higher multiplicity. In other words, Nature is likely to realize them in more ways. If W_1 and W_2 are two distributions, with corresponding entropies of H_1 and H_2 , then the ratio W_2/W_1 is the relative preference of W_2 over W_1 . Since $W_2/W_1 \sim \exp[N(H_2 - H_1)]$, when N becomes large (such as the Avogadro number), the relative preference "becomes so overwhelming that exceptions to it are never seen; and we call it the Second Law of Thermodynamics." (Jaynes, 1982).

The problem may now be expressed in terms of constrained optimization as follows:

$$\underset{\{N_k\}}{\text{Maximize}} \quad \log W = -N \sum_{k=1}^s \left(\frac{N_k}{N} \right) \log \left(\frac{N_k}{N} \right)$$

subject to

$$\sum_{k=1}^s N_k E_k = E, \quad \text{and}$$

$$\sum_{k=1}^s N_k = N.$$

The solution yields surprisingly rich results which would not be attainable even if the individual trajectories of all the molecules in the closed spaces were calculated. The efficiency of the method reveals that in fact, such voluminous calculations would have canceled each other out, and were actually irrelevant to the problem. A similar idea is seen in the chapter on genetic algorithms, where ignorance can be seemingly

exploited and irrelevant information, even if assumed, would be eliminated from the solution.

The technique has been used in artificial intelligence (see for example, [Lippman, 1988; Jaynes, 1991; Kane, 1991]), and in solving problems in business and economics (see for example, [Jaynes, 1991; Grandy, 1991; Zellner, 1991]).

4.2 Examples

We will see how the principle is used in solving problems involving some type of prior information which is used as a constraint on the problem. For simplicity, we will deal with problems involving one random variable Θ having n values, and call the associated probabilities p_i . For all the problems, the goal is to choose a probability distribution from among many possible ones which has the maximum entropy.

No prior information whatsoever. The problem may be formulated using the Lagrange multiplier λ for the single constraint as:

$$\underset{\{p_i\}}{\text{Max}} \quad g(\{p_i\}) = - \sum_{i=1}^n p_i \ln p_i + \lambda \left[\sum_{i=1}^n p_i - 1 \right].$$

The solution is obtained as follows: Hence, $p_i = 1/n$, $i = 1, \dots, n$ is the MaxEnt assignment, which confirms the intuition on the non-informative prior.

Suppose the expected value of Θ is μ_0 . We have two constraints in this problem: the first is the usual constraint on the probabilities summing to one; the second is the given information - expected value of Θ is μ_0 . We use the Lagrange multipliers λ_1 and λ_2 for the two constraints respectively. The problem statement follows:

$$\frac{\partial g}{\partial p_i} = -1 - \ln p_i + \lambda = 0$$

$$\Rightarrow \ln p_i = \lambda - 1$$

$$\Rightarrow p_i = e^{\lambda-1} \quad \forall i = 1, \dots, n,$$

$$\frac{\partial g}{\partial \lambda} = \sum_{i=1}^n p_i = 1$$

$$\Rightarrow \sum_{i=1}^n e^{\lambda-1} = 1$$

$$\Rightarrow n e^{\lambda-1} = 1$$

$$\Rightarrow n p_i = 1$$

$$\Rightarrow p_i = \frac{1}{n} \quad \forall i = 1, \dots, n.$$

$$\text{Max}_{\{p_i\}} g(\{p_i\}) = - \sum_{i=1}^n p_i \ln p_i + \lambda_1 \left[\sum_{i=1}^n p_i - 1 \right] + \lambda_2 \left[\sum_{i=1}^n (\theta_i p_i - \mu_\theta) \right].$$

This can be solved in the usual way by taking partial derivatives of $g()$ w.r.t. p_i , λ_1 , and

λ_2 , and equating them to zero. We obtain:

$$p_i = e^{\lambda_1 + \lambda_2 \theta_i - 1}, \quad \text{and}$$

$$\sum_{i=1}^n \theta_i e^{\lambda_2 \theta_i} = \mu_\theta \sum_{i=1}^n e^{\lambda_2 \theta_i}.$$

Writing

$$x = e^{\lambda_2},$$

we get

$$\begin{aligned}\sum_{i=1}^n \theta_i x^{\theta_i} &= \mu_\theta \sum_{i=1}^n x^{\theta_i} \\ \Leftrightarrow \sum_{i=1}^n (\theta_i - \mu_\theta) x^{\theta_i} &= 0\end{aligned}$$

which is a polynomial in x , whose roots can be determined numerically.

For example, let $n = 3$, θ take values $\{1, 2, 3\}$, $\mu_\theta = 1.25$. Solving as above and taking the appropriate roots, we obtain

$$\lambda_1 \approx 2.2752509, \lambda_2 \approx -1.5132312, \text{ giving}$$

$$p_1 \approx 0.7882, p_2 \approx 0.1671, \text{ and } p_3 \approx 0.0382.$$

Partial knowledge of probabilities. Suppose we know p_i , $i = 1, \dots, k$. Since we have $n-1$ degrees of freedom in choosing p_i , assume $k \leq n-2$ to make the example non-trivial. Then, the problem may be formulated as:

$$\max_{\{p_i\}} g(\{p_i\}) = - \sum_{i=k+1}^n p_i \ln p_i + \lambda \left[\sum_{i=k+1}^n p_i + q - 1 \right],$$

$$\text{where } q = \sum_{i=1}^k p_i.$$

Solving, we obtain

$$p_i = \frac{1-q}{n-k}, \quad \forall i = k+1, \dots, n.$$

This is again fairly intuitive: the remaining probability $1-q$ is distributed non-informatively over the rest of the probability space. For example, if $n = 4$, $p_1 = 0.5$, and $p_2 = 0.3$, then $k = 2$, $q = 0.8$, and $p_3 = p_4 = (1 - 0.8)/(4 - 2) = 0.2/2 = 0.1$.

Note that the first case is a special case of the last one, with $q = k = 0$.

The technique can be extended to cover prior knowledge expressed in the form of probabilistic knowledge bases by using two key MaxEnt solutions: non-informativeness (as covered in the last example above), and statistical independence of two random variables given no knowledge to the contrary (in other words, given two probability distributions f and g over two random variables X and Y respectively, and no further information, the MaxEnt joint probability distribution h over $X*Y$ is obtained as $h = f*g$).

CHAPTER 5 THE PRINCIPAL-AGENT PROBLEM

5.1 Introduction

5.1.1 The Agency Relationship

The principal-agent problem arises in the context of the agency relationship in social interaction. The agency relationship occurs when one party, the agent, contracts to act as a representative of another party, the principal, in a particular domain of decision problems.

The principal-agent problem is a special case of a dynamic two-person game. The principal has available to her a set of possible compensation schemes, out of which she must select one that both motivates the agent and maximizes her welfare. The agent also must choose a compensation scheme which maximizes his welfare, and he does so by accepting or rejecting the compensation schemes presented to him by the principal. Each compensation package he considers implicitly influences him to choose a particular (possibly complex) action or level of effort. Every action has associated with it certain disutilities to the agent, in that he must expend a certain amount of effort and/or expense. It is reasonable to assume that the agent will reject outright any compensation package which yields less than that which can be obtained elsewhere in the market. This assumption is in turn based on the assumptions that the agent is knowledgeable about his

"reservation constraint", and that he is free to act in a rational manner. The assumption of rationality also applies to the principal. After agreeing to a contract, the agent proceeds to act on behalf of the principal, which in due course yields a certain outcome. The outcome is not only dependent on the agent's actions but also on exogenous factors. Finally the outcome, when expressed in monetary terms, is shared between the principal and the agent in the manner decided upon by the selected compensation plan.

The specific ways in which the agency relationship differs from the usual employer-employee relationship are (Simon, 1951):

- (1) The agent does not recognize the authority of the principal over specific tasks the agent must do to realize the output.
- (2) The agent does not inform the principal about his "area of acceptance" of desirable work behavior.
- (3) The work behavior of the agent is not directly (or costlessly) observable by the principal.

Some of the first contributions to the analysis of principal-agent problems can be found in Simon (1951), Alchian & Demsetz (1972), Ross (1973), Sitglitz (1974), Jensen & Meckling (1976), Shavell (1979a, 1979b), Holmstrom (1979, 1982), Grossman & Hart (1983), Rees (1985), Pratt & Zeckhauser (1985), and Arrow (1986).

There are three critical components in the principal-agent model: the technology, the informational assumptions, and the timing. Each of these three components is described below.

5.1.2 The Technology Component of Agency

The technology component deals with the type and number of variables involved (for example, production variables, technology parameters, factor prices, etc.), the type and the nature of functions defined on these variables (for example, the type of utility functions, the presence of uncertainty and hence the existence of probability distribution functions, continuity, differentiability, boundedness, etc.), the objective function and the type of optimization (maximization or minimization), the decision criteria on which optimization is carried out (expected utility, weighted welfare measures, etc.), the nature of the constraints, and so on.

5.1.3 The Information Component of Agency

The information component deals with the private information sources of the principal and the agent, and information which is public (i.e. known to both the parties and costlessly verifiable by a third party, such as a court). This component of the model addresses the question, "who knows what?". The role of the informational assumption in agency is as follows:

- (a) it determines how the parties act and make decisions (such as offer payment schemes or choose effort levels),
- (b) it makes it possible to identify or design communication structures,
- (c) it determines what additional information is necessary or desirable for improved decision making, and

- (d) it enables the computation of the cost of maintaining or establishing communication structures, or the cost of obtaining additional information.

For example, one usual assumption in the principal-agent literature is that the agent's reservation level is known to both parties. As another example of the way in which additional information affects the decisions of the principal, note that the principal, in choosing a set of compensation schemes for presenting to the agent, wishes to maximize her welfare. It is in her interest, therefore, to make the agent accept a payment scheme which induces him to choose an effort level that will yield a desired level of output (taking into consideration exogenous risk). The principal would be greatly assisted in her decision making if she had knowledge of the "function" which induces the agent to choose an effort level based on the compensation scheme, and also knowledge of the hidden characteristics of the agent such as his utility of income, disutility of effort, risk attitude, reservation constraint, etc. Similarly, the agent would be able to take better decisions if he were more aware of his risk attitude, disutility of effort and exogenous factors. Any information, even if imperfect, would reduce either the magnitude or the variance of risk or both. However, better information for the agent does not always imply that the agent will choose an act or effort level that is also optimal for the principal. In some cases, the total welfare of the agency may be reduced as a result (Christensen, 1981).

The gap in information may be reduced by employing a system of messages from the agent to the principal. This system of messages may be termed a "communication structure" (Christensen, 1981). The agent chooses his action by observing a signal from

his private information system after he accepts a particular compensation scheme from the principal subject to its satisfying the reservation constraint. This signal is caused by the combination of the compensation scheme, an estimate of exogenous risk by the agent based on his prior information or experience, and the agent's knowledge of his risk attitude and disutility of action. The communication structure agreed upon by both the principal and the agent allows the agent to send a message to the principal. It is to be noted that the agency contract can be made contingent on the message, which is jointly observable by both the parties. The compensation scheme considers the message(s) as one (some) of the factors in the computation of the payment to the agent, the other of course being the output caused by the agent's action. Usually, formal communication is not essential, as the principal can just offer the agent a menu of compensation schemes, and allow the agent to choose one element of the menu.

5.1.4 The Timing Component of Agency

Timing deals with the sequence of actions taken by the principal and the agent, and the time when they commit themselves to specific decisions (for example, the agent may choose an effort level before or after observing some signal about exogenous risk).

Below is one example of timing (T denotes time):

- T1. The principal selects a particular compensation scheme from a set of possible compensation schemes.
- T2. The agent accepts or rejects the suggested compensation scheme depending on whether it satisfies his reservation constraint or not.

- T3. The agent chooses an action or effort level from a set of possible actions or effort levels.
- T4. The outcome occurs as a function of the agent's actions and exogenous factors which are unknown or known only with uncertainty.

Another example of timing is when a communication structure with signals and messages is involved (Christensen, 1981):

- T1. The principal designs a compensation scheme.
- T2. Formation of the agency contract.
- T3. The agent observes a signal.
- T4. The agent chooses an act and sends a message to the principal.
- T5. The output occurs from the agent's act and exogenous factors.

Variations in the principal-agent problems are caused by changes in one or more of these components. For example, some principal-agent problems are characterized by the fact that the agent may not be able to enforce the payment commitments of the principal. This situation occurs in some of the relationships in the context of regulation. Another is the possibility of renegotiation or review of the contract at some future date.

Agency theory, dealing with the above market structure, gives rise to a variety of problems caused by the presence of factors such as the influence of externalities, limited observability, asymmetric information, and uncertainty (Gjesdal, 1982).

5.1.5 Limited Observability, Moral Hazard, and Monitoring

An important characteristic of principal-agent problems - limited observability of the agent's actions - gives rise to moral hazard. Moral hazard is a situation in which one party (say, the agent) may take actions detrimental to the principal and which cannot be perfectly and/or costlessly observed by the principal (see for example, [Holmstrom, 1979]). Formally, perfect observation might very well impose "infinite" costs on the principal. The problem of unobservability is usually addressed by designing monitoring systems or signals which act as estimators of the agent's effort. The selection of monitoring signals and their value is discussed for the case of costless signals in Harris and Raviv (1979), Holmstrom (1979), Shavell (1979), Gjesdal (1982), Singh (1985), and Blickle (1987). Costly signals are discussed for three cases in Blickle (1987).

On determining the appropriate monitoring signals, the principal invites the agent to select a compensation scheme from a class of compensation schemes which she, the principal, compiles. Suppose the principal determines monitoring signals s_1, \dots, s_n , and has a compensation scheme $c(q, s_1, \dots, s_n)$, where q is the output, which the agent accepts. There is no agreement between the principal and the agent as to the level of the effort e . Since the signals $s_i, i = 1, \dots, n$ determine the payoff and the effort level e of the agent (assuming the signals have been chosen carefully), the agent is thereby induced to an effort level which maximizes the expected utility of his payoff (or some other decision criterion). The only decision still in the agent's control is the choice of how much payoff he wants; the assumption is that the agent is rational in an economic sense. The principal's residuum is the output q less the compensation $c(\cdot)$. The principal

structures the compensation scheme $c(\cdot)$ in such a way as to maximize the expected utility of her residuum (or some other decision criterion). In this manner, the principal induces desirable work behavior in the agent.

It has been observed that "the source of moral hazard is not unobservability but the fact that the contract cannot be conditioned on effort. Effort is **noncontractible**." (Rasmusen, 1989). This is true when the principal observes shirking on the part of the agent but is unable to prove it in a court of law. However, this only implies that a contract on effort is imperfectly enforceable. Moral hazard may be alleviated in cases where effort is contracted, and where both limited observability and a positive probability of proving non-compliance exist.

5.1.6 Informational Asymmetry, Adverse Selection, and Screening

Adverse selection arises in the presence of informational asymmetry which causes the two parties to act on different sets of information. When perfect sharing of information is present and certain other conditions are satisfied, first-best solutions are feasible (Sappington and Stiglitz, 1987). Typically however, adverse-selection exists.

While the effect of moral hazard makes itself felt when the agent is taking actions (say, production or sales), adverse selection affects the formation of the relationship, and may give rise to inefficient (in the second-best sense) contracts. In the information-theoretic approach, we can think of both being caused by lack of information. This is variously referred to as the dissimilarity between private information systems of the agent

and the firm, or the unobservability or ignorance of "hidden characteristics" (in the latter sense, moral hazard is caused by "hidden effort or actions").

In the theory of agency, the hidden characteristic problem is addressed by designing various sorting and screening mechanisms, or communication systems that pass signals or messages about the hidden characteristics (of course, the latter can also be used to solve the moral hazard problem).

On the one hand, the screening mechanisms can be so arranged as to induce the target party to select by itself one of the several alternative contracts (or "packages"). The selection would then reveal some particular hidden characteristic of the party. In such cases, these mechanisms are called "self-selection" devices. See, for example, Spremann (1987) for a discussion of self-selection contracts designed to reveal the agent's risk attitude. On the other hand, the screening mechanisms may be used as indirect estimators of the hidden characteristics, as when aptitude tests and interviews are used to select agents.

The significance of the problem caused by the asymmetry of information is related to the degree of lack of trust between the parties to the agency contract which, however, may be compensated for by observation of effort. However, most real life situations involving an agency relationship of any complexity are characterized not only by a lack of trust but also by a lack of observability of the agent's effort. The full context to the concept of information asymmetry is the fact that each party in the agency relationship is either unaware or has only imperfect knowledge of certain factors which are better known to the other party.

5.1.7 Efficiency of Cooperation and Incentive Compatibility

In the absence of asymmetry of information, both principal and agent would cooperatively determine both the payoff and the effort or work behavior of the agent. Subsequently, the "game" would be played cooperatively between the principal and the agent. This would lead to an efficient agreement termed the first-best design of cooperation. First-best solutions are often absent not merely because of the presence of externalities but mainly because of adverse selection and moral hazard (Spremann, 1987).

Let $F = \{ (c,e) \}$, where compensation c and effort e satisfy the principal's and the agent's decision criteria respectively. In other words, F is the set of first-best designs of cooperation, also called efficient designs with respect to the principal-agent decision criteria. Now, suppose that the agent's action e is induced as above by a function I : $I(c) = e$. Let $S = \{ (c,I(c)) \}$ -- i.e. S denotes the set of designs feasible under information asymmetry. If it were not the case that $F \cap S = \emptyset$, then efficient designs of cooperation would be easily induced by the principal. Situations where this occurs are said to be incentive compatible. In all other cases, the principal has available to her only second-best designs of cooperation, which are defined as those schemes that arise in the presence of information asymmetry.

5.1.8 Agency Costs

There are three types of agency costs (Schneider, 1987):

- (1) the cost of monitoring the hidden effort of the agent,
- (2) the bonding costs of the agent, and

- (3) the residual loss, defined as the monetary equivalent of the loss in welfare of the principal caused by the actions taken by the agent which are non-optimal with respect to the principal.

Agency costs may be interpreted in the following two ways:

- (1) they may be used to measure the "distance" between the first-best and the second-best designs;
- (2) they may be looked upon as the value of information necessary to achieve second-best designs which are arbitrarily close to the first-best designs.

Obviously, the value of perfect information should be considered as an upper bound on the agency costs (see for example, [Jensen and Meckling, 1976]).

5.2 Formulation of the Principal-Agent Problem

The following notation and definitions will be used throughout:

D: the set of decision criteria, such as {maximin, minimax, maximax, minimin, minimax regret, expected value, expected loss,...}. We use $\Delta \in D$.

Δ_P : the decision criterion of the principal.

Δ_A : the decision criterion of the agent.

U_P : the principal's utility function.

U_A : the agent's utility function.

C: the set of all compensation schemes. We use $c \in C$.

E: the set of actions or effort levels of the agent. We use $e \in E$.

Θ : a random variable denoting the true state of nature.

Θ_P : a random variable denoting the principal's estimate of the state of nature.

Θ_A : a random variable denoting the agent's estimate of the state of nature.

q : output realized from the agent's actions (and possibly the state of nature).

q_P : monetary equivalent of the principal's residuum. Note that $q_P = q - c(\cdot)$, where c may depend on the output and possibly other variables.

Output/outcome. The goal or purpose of the agency relationship, such as sales, services or production, is called the output or the outcome.

Public knowledge/information. Knowledge or information known to both the principal and the agent, and also a third enforcement party, is termed public knowledge or information. A contract in agency can be based only on public knowledge (i.e. observable output or signals).

Private knowledge/information. Knowledge or information known to either the principal or the agent but not both is termed private knowledge or information.

State of nature. Any events, happenings, occurrences or information which are not in the control of the principal or the agent and which affect the output of the agency directly through the technology constitute the state of nature.

Compensation. The economic incentive to the agent to induce him to participate in the agency is called the compensation. This is also called wage, payment or reward.

Compensation scheme. The package of benefits and output sharing rules or functions that provide compensation to the agent is called the compensation scheme. Also called contract, payment function or compensation function.

The word "scheme" is used here instead of "function" since complicated compensation packages will be considered as an extension later on. In the literature, the word "scheme" may be seen, but it is used in the sense of "function", and several nice properties are assumed for the function (such as continuity, differentiability, and so on). Depending on the contract, the compensation may be negative - a penalty for the agent. Typical components of the compensation functions considered in the literature are rent (fixed and possibly negative), and share of the output.

The principal's residuum. The economic incentive to the principal to engage in the agency is the principal's residuum. The residuum is the output (expressed in monetary terms) less the compensation to the agent. Hence, the principal is sometimes called the residual claimant.

Payoff. Both the agent's compensation and the principal's residuum are called the payoffs.

Reservation welfare (of the agent). The monetary equivalent of the best of the alternative opportunities (with other competing principals, if any) available to the agent is known as the reservation welfare of the agent. Accordingly, it is the minimum compensation that induces an agent to accept the contract, but not necessarily induce him to his best effort level. Also known as reservation utility or individual utility, it is variously denoted in the literature as m or \bar{U} .

Disutility of effort. The cost of inputs which the agent must supply himself when he expends effort contributes to disutility, and hence is called the disutility of effort.

Individual rationality constraint (IRC). The agent's (expected) utility of net compensation (compensation from the principal less his disutility of effort) must be at least as high as his reservation welfare. This constraint is also called the participation constraint.

When a contract violates the individual rationality constraint, the agent rejects it and prefers unemployment instead. Such a contract is not necessarily "bad", since different individuals have different levels of reservation welfare. For example, financially independent individuals may have higher than usual reservation welfare levels, and might very well prefer leisure to work even when contracts are attractive to most other people.

Incentive compatibility constraint (ICC). A contract will be acceptable to the agent if it satisfies his decision criterion on compensation, such as maximization of expected utility of net compensation. This constraint is called the incentive compatibility constraint.

Development of the problem: Model 1. We develop the problem from simple cases involving the least possible assumptions on the technology and informational constraints, to those having sophisticated assumptions. Corresponding models from the literature are reviewed briefly in section 1.3.

A. Technology:

- (a) fixed compensation, $C \equiv$ set of fixed compensations, $\bar{U} \in C$;

- (b) output $q \equiv q(e)$; assume $q(0) = 0$;
- (c) existence of nonseparable utility functions;
- (d) decision criterion: maximization of utility;
- (e) no uncertainty in the state of nature.

B. Public information:

- (a) compensation scheme, c ;
- (b) range of possible outputs, Q ;
- (c) \bar{U} .

Information private to the principal: U_p

Information private to the agent:

- (a) U_A ;
- (b) disutility of effort, d ;
- (c) range of effort levels, e .

C. Timing:

- (1) the principal makes an offer of fixed wage c ;
- (2) the agent either rejects or accepts the offer;
- (3) if he accepts it, exerts effort level e ;
- (4) output $q(e)$ results;

(5) sharing of output according to contract.

D. Payoffs:

Case 1: Agent rejects contract, i.e. $e = 0$;

$$\pi_p = U_p[q(e)] = U_p[q(0)] = U_p[0].$$

$$\pi_A = U_A[\bar{U}].$$

Case 2: Agent accepts contract;

$$\pi_p = U_p[q(e) - c].$$

$$\pi_A = U_A[c - d(e)].$$

E. The principal's problem:

$$(M1.P1) \quad \max_{c \in C} \max_{q \in Q} U_p[q - c]$$

such that

$$c \geq \bar{U}. \quad (IRC)$$

Suppose $C^* \subseteq C$ is the solution set of M1.P1. The principal picks $c^* \in C^*$ and offers it to the agent.

The agent's problem:

(M1.A1) For a given c^* ,

$$\max_{e \in E} U_A[c^* - d(e)].$$

Suppose $E^* \subseteq E$ is the solution set of M1.A1. The agent selects $e^* \in E^*$.

F. The solution:

- (a) the principal offers $c^* \in C^*$ to the agent;
- (b) the agent accepts the contract;
- (c) the agent exerts effort $e^*(c^*) \in E^*$;
- (d) output $q(e^*(c^*))$ occurs;
- (e) payoffs:

$$\pi_P = U_P[q(e^*(c^*)) - c^*];$$

$$\pi_A = U_A[c^* - d(e^*(c^*))].$$

Notes:

1. The agent accepts the contract in F.b since IRC is present in M1.P1, and C^* is nonempty since $\bar{U} \in C$.
2. Effort of the agent is a function of the offered compensation.
3. Since one of the informational assumptions was that the principal does not know the agent's utility function, \bar{U} is a compensation rather than the agent's utility of compensation, so $U_A(\bar{U})$ is meaningful.

G. Variations:

1. The principal offers C^* to the agent instead of a $c^* \in C^*$. The agent's problem then becomes:

$$(M1.A2) \text{ Max}_{c^* \in C^*} \max_{e \in E} U_A[c^* - d(e)].$$

The first three steps in the solution then become:

- (a) the principal offers C^* to the agent;

- (b) the agent accepts the contract;
 - (c) the agent picks an effort level e^* which is a solution to M1.A2 and reports the corresponding c^* (or its index if appropriate) to the principal.
2. The agent may decide to solve an additional problem: from among two or more competing optimal effort levels, he may wish to select a minimum effort level. Then, his problem would be:

$$(M1.A3) \text{ Min}_{e^*} d(e^*)$$

such that

$$e^* \in \operatorname{argmax}_{e \in E} U_A[c^* - d(e)].$$

Example:

Let $E = \{e_1, e_2, e_3\}$,

$$C^* = \{c_1, c_2, c_3\}.$$

Suppose,

$$c_1(q(e_1)) = 5, d(e_1) = 2;$$

$$c_2(q(e_2)) = 6, d(e_2) = 3;$$

$$c_3(q(e_3)) = 6, d(e_3) = 4;$$

The net compensation to the agent in choosing the three effort levels is 3, 3, and 2 respectively. Assuming $d(e)$ is monotone increasing in e , the agent chooses e_1 to e_2 , and so prefers compensation c_1 to c_2 .

3. We assumed \bar{U} is public knowledge. If this were not so, then the agent has to test all offers to see if they are at least as high as the utility of his reservation welfare. The two problems then become:

$$(M1.P2) \quad \text{Max}_{c \in C} \max_{q \in Q} U_P[q - c]$$

and

$$(M1.A4) \quad \text{Max}_{e \in E} U_A[c^* - d(e)]$$

such that

$$c^* \geq U_A[\bar{U}], \quad (\text{IRC})$$

$$c^* \in \text{argmax } M1.P2.$$

In this case, there is a distinct possibility of the agent rejecting an offer of the principal.

4. Note that in most realistic situations, a distinction must be made between the reservation welfare and the agent's utility of the reservation welfare. Otherwise, merely using IRC with the reservation welfare in M1.P1 may not satisfy the agent's constraint. On the other hand, $\bar{U} \equiv U_A(\bar{U})$ implies knowledge of U_A by the principal, a complication which yields a completely different model.

When $\bar{U} \neq U_A(\bar{U})$, the following two problems occur:

$$(M1.P3) \quad \text{Max}_{c \in C} \max_{q \in Q} U_P(q - c)$$

such that

$$c \geq \bar{U}.$$

$$(M1.A5) \quad \text{Max}_{e \in E} U_A(c^* - d(e))$$

such that

$$c^* \geq U_A(\bar{U}), \quad (\text{IRC})$$

$$c^* \in \operatorname{argmax} M1.P3.$$

In other words, the principal solves her problem the best way she can, and hopes the solution is acceptable to the agent.

5. Negotiation. Negotiation of a contract can occur in two contexts:

- (a) when there is no solution to the initial problem, the agent may communicate to the principal his reservation welfare, and the principal may design new compensation schemes or revise her old schemes so that a solution may be found. This type of negotiation also occurs in the case of problems M1.P3 and M1.A5.
- (b) The principal may offer $c^* \in \operatorname{argmax}_{c \in C} M1.P1$. The agent either accepts it or does not; if he does not, then the principal may offer another optimal contract, if any. This interaction may continue until either the agent accepts some compensation scheme or the principal runs out of optimal compensations.

Development of the problem: Model 2. This model differs from the first by incorporating uncertainty in the state of nature, and conditioning the compensation functions on the output.

A. Technology:

- (a) presence of uncertainty in the state of nature;
- (b) compensation scheme $c \equiv c(q)$;
- (c) output $q \equiv q(e, \theta)$;
- (d) existence of known utility functions for the agent and the principal;
- (e) disutility of effort for the agent is monotone increasing in effort e ;

B. Public information:

- (a) presence of uncertainty, and range of Θ ;
- (b) output function q ;
- (c) payment functions c ;
- (d) range of effort levels of the agent.

Information private to the principal:

- (a) the principal's utility function;
- (b) the principal's estimate of the state of nature.

Information private to the agent:

- (a) the agent's utility function;
- (b) the agent's estimate of the state of nature;
- (c) disutility of effort;
- (d) reservation welfare;

C. Timing:

- (a) the principal determines the set of all compensation schemes that maximize her expected utility;
- (b) the principal presents this set to the agent as the set of offered contracts;
- (c) the agent picks from this set of compensation schemes a compensation scheme that maximizes his net compensation, and a corresponding effort level;
- (d) a state of nature occurs;
- (e) an output results;
- (f) sharing of the output takes place as contracted.

D. Payoffs:

Case 1: Agent rejects contract, i.e. $e = 0$;

$$\pi_P = U_P[q(e, \theta)] = U_P[q(0, \theta)].$$

$$\pi_A = U_A[\bar{U}].$$

Case 2: Agent accepts contract;

$$\pi_P = U_P[q(e, \theta) - c(q)].$$

$$\pi_A = U_A[c(q) - d(e)].$$

E. The principal's problem:

$$(M2.P) \quad \text{Max}_{c \in C} \text{Max}_{\theta \in E} E^{\theta_P} [U_P[q(e, \theta) - c(q(e, \theta))]]$$

where the expectation $E(\cdot)$ is given by (assuming the usual regularity conditions)

$$\int_0^{\bar{\theta}} U_P[q(e, \theta) - c(q(e, \theta))] f_{\theta_P}(\theta) d\theta$$

where

$$\theta \in [\underline{\theta}, \bar{\theta}], \quad \text{and}$$

$f(\theta)$ is the distribution assigned by the principal.

The agent's problem:

$$(M2.A) \quad \text{Max}_{c \in C} \text{Max}_{e \in E} E^{\theta_A} [U_A[c(q(e, \theta)) - d(e)]]$$

subject to

$$E^{\theta_A} [c(q(e, \theta)) - d(e)] \geq \bar{U}, \quad (IRC)$$

$$c \in \text{argmax}(M2.P).$$

where the expectation $E(\cdot)$ is given as usual by

$$\int_0^{\bar{\theta}} U_A[q(e, \theta) - c(q(e, \theta))] f_{\theta_A}(\theta) d\theta.$$

F. The solution:

- (a) The agent selects $c^* \in C^*$, and a corresponding effort e^* which is a solution to M2.A;
- (b) a state of nature θ occurs;
- (c) output $q(e^*, \theta)$ is generated;
- (d) payoffs:

$$\pi_P = U_P[q(e^*, \theta) - c^*(q(e^*, \theta))];$$

$$\pi_A = U_A[c^*(q(e^*, \theta)) - d(e^*)].$$

Development of the problem: Model 3. In this model, the strongest possible assumption is made about information available to the principal: the principal has complete knowledge of the utility function of the agent, his disutility of effort, and his reservation welfare. Accordingly, the principal is able to make an offer of compensation which satisfies the decision criterion of the agent and his constraints. In other words, the two problems are treated as one. The assumptions are as in model 2, so only the statement of the problem will be given below.

The problem:

$$\text{Max}_{c \in C, e^* \in E} E U_P[q(e^*, \theta) - c(q(e^*, \theta))]$$

subject to

$$E U_A[c(q(e^*, \theta)) - d(e^*)] \geq \bar{U}, \quad (IRC)$$

$$e^* \in \text{argmax}_{e \in E, c \in C} \{E U_A[c(q(e, \theta)) - d(e)]\}. \quad (ICC)$$

5.3 Main Results in the Literature

Several results from basic agency models will be presented using the framework established in the development of the problem. The following will be presented for each model:

Technology,

Information,

Timing,

Payoffs, and

Results.

It must be noted that the literature rarely presents such an explicit format; rather, several assumptions are often buried within the results, or implied or just not stated. Only by trying an algorithmic formulation is it possible to unearth unspecified assumptions. In many cases, some of the factors are assumed for the sake of formal completeness, even though the original paper neither mentions nor uses those factors in its results. This type of modeling is essential when the algorithms are implemented subsequently using a knowledge-intensive methodology.

One recurrent example of incomplete specification is the treatment of the agent's individual rationality constraint (IRC). The principal has to pick a compensation which satisfies IRC. However, some consistency in using IRC is necessary. The agent's reservation welfare \bar{U} is also a compensation (albeit a default one). The agent must

check one of two constraints to verify that the offered compensation indeed meets his reservation welfare:

$$c \geq \bar{U} \quad \text{or} \quad U_A(c) \geq U_A(\bar{U}).$$

If the principal picks a compensation which satisfies $c \geq \bar{U}$, it is not necessary that $U_A(c) \geq U_A(\bar{U})$ be also satisfied. However, using $U_A(c) \geq \bar{U}$ for the IRC, where \bar{U} is treated "as if" it were $U_A(\bar{U})$, implies knowledge of the agent's utility on the part of the principal.

The difference between the two situations is of enormous significance if the purpose of analysis is to devise solutions to real-world problems. In the literature, this distinction is conveniently overlooked. If all such vagueness in the technological, informational and temporal assumptions was to be systematically eliminated, the analysis might change in a way not intended in the original literature. Hence, the main results in the literature will be presented as they are.

5.3.1 Model 1: The Linear-Exponential-Normal Model

This name of the model (Spremann, 1987) derives from the nature of three crucial parameters: the payoff functions are linear, the utility functions are exponential, and the exogenous risk has a normal distribution. Below is a full description.

Technology:

- (a) compensation is the sum of a fixed rent r and a share s of the output q : $c(q)$
- $$\equiv r + sq;$$

- (b) presence of uncertainty in the state of nature, denoted by Θ , where $\Theta \sim N(0, \sigma^2)$;
- (c) the set of effort levels of the agent, $E = [0, \frac{1}{2}]$; effort is induced by compensation;
- (d) output $q \equiv q(e, \theta) \equiv e + \theta$;
- (e) the agent's disutility of effort is $d \equiv d(e) \equiv e^2$;
- (f) the principal's utility U_P is linear (the principal is risk neutral);
- (g) the agent has constant risk aversion $\alpha > 0$, and his utility is $U_A(w) = -\exp(-\alpha w)$, where w is his net compensation (also called the wealth);
- (h) the certainty equivalent of wealth, denoted V , is defined as:

$$V(w) \equiv U^{-1}[E^\theta(U(w))]$$
, where U denotes the utility function, E^θ is the expectation with respect to Θ ; as usual, subscripts P or A on V denote the principal or the agent respectively;
- (i) the decision criterion is maximization of expected utility.

Public information:

- (a) compensation scheme $c(q; r, s)$;
- (b) output q ;
- (c) distribution of Θ ;
- (d) agent's reservation welfare \bar{U} ;
- (e) agent's risk aversion α .

Information private to the principal:

Utility of residuum, U_p .

Information private to the agent:

- (a) selection of effort given the compensation;
- (b) utility of welfare;
- (c) disutility of effort.

Timing:

- (a) the principal offers a contract (r, s) to the agent;
- (b) the agent's effort e is induced by the compensation scheme;
- (c) a state of nature occurs;
- (d) the agent's effort and the state of nature give rise to output;
- (e) sharing of the output takes place.

Payoffs:

$$\begin{aligned}
 \pi_p &= U_p[q - (r + sq)] \\
 &= U_p[e(r, s) + \theta_0 - (r + s(e(r, s) + \theta_0))] \\
 \pi_A &= U_A[r + sq - d(e(r, s))] \\
 &= U_A[r + s(e(r, s) + \theta_0) - d(e(r, s))],
 \end{aligned}$$

where $e(r, s)$ is the function which induces effort based on compensation, and θ_0 is the realized state of nature.

Results:

Result 1.1: The optimal effort level of the agent given a compensation scheme (r,s) is denoted e^* , and is obtained by straightforward maximization to yield:

$$e^* \equiv e^*(r,s) = s/2.$$

This shows that the rent r and the reservation welfare \bar{U} have no impact on the selection of the agent's effort.

Result 1.2: A necessary and sufficient condition for IRC to be satisfied for a given compensation scheme (r,s) is:

$$r \geq \bar{U} - \frac{s^2(1 - 2\alpha\sigma^2)}{4}$$

Result 1.3: The optimal compensation scheme for the principal is $c^* \equiv (r^*, s^*)$, where

$$s^* = \frac{1}{1 + 2\alpha\sigma^2}, \quad \text{and}$$

$$r^* = \bar{U} - \frac{1 - 2\alpha\sigma^2}{4s^{*2}}.$$

Corollary 1.3: The agent's optimal effort given a compensation scheme (r^*, s^*) is (using result 1.1):

$$e^* = \frac{1}{2(1 + 2\alpha\sigma^2)}$$

Result 1.4: Suppose $2\alpha\sigma^2 > 1$. Then, an increase in share s requires an increase in rent r (in order to satisfy IRC).

To see this, suppose we increase the share s by δ ,

$s_0 = s + \delta$, $0 < \delta \leq 1-s$. From Result 1.2, for IRC to hold we need,

$$\begin{aligned}
 r_0 &\geq \bar{U} - \frac{s_0^2(1 - 2\alpha\sigma^2)}{4} \\
 &= \bar{U} - \frac{(s + \delta)^2(1 - 2\alpha\sigma^2)}{4} \\
 &= \bar{U} - \frac{(1 - 2\alpha\sigma^2)[s^2 + 2s\delta + \delta^2]}{4} \\
 &= \bar{U} - \frac{(1 - 2\alpha\sigma^2)}{4}s^2 - \frac{(2s\delta + \delta^2)(1 - 2\alpha\sigma^2)}{4} \\
 &= r - \frac{(2s\delta + \delta^2)(1 - 2\alpha\sigma^2)}{4} \\
 &\geq r \quad (\because 1 < 2\alpha\sigma^2).
 \end{aligned}$$

Result 1.5: The welfare attained by the agent is \bar{U} , while the principal's welfare is given by:

$$\frac{1}{4s^*} - \bar{U}.$$

Result 1.6: The principal prefers agents with lower risk aversion. This is immediate from the fact that the principal's welfare is decreasing in the agent's risk aversion for a given σ^2 and \bar{U} .

Result 1.7: Fixed fee arrangements are non-optimal, no matter how large the agent's risk aversion. This is immediate from the fact that

$$s^* = \frac{1}{1 + 2\alpha\sigma^2} > 0 \quad \forall \alpha > 0.$$

Result 1.8: It is the connection between unobservability of the agent's effort and his risk aversion that excludes first-best solutions.

5.3.2 Model 2

This model (Gjesdal, 1982) deals with two problems:

- (a) choosing an information system, and
- (b) designing a sharing rule based on the information system.

Technology:

- (a) presence of uncertainty, Θ ;
- (b) finite effort set of the agent; effort has several components, and is hence treated as a vector;
- (c) output q is a function of the agent's effort and the state of nature θ ; the range of output levels is finite;

- (d) presence of a finite number of public signals;
- (e) presence of a set of public information systems (i.e. signals), including non-informative and randomized systems, the output being treated as one of the informative information systems;
- (f) costlessness of public information systems;
- (g) compensation schemes are based on signals about effort or output or both.

Public information:

- (a) distribution of the state of nature, Θ ;
- (b) output levels;
- (c) common information systems which are non-informative and randomizing;
- (d) U_A .

Information private to the principal: utility function, U_P .

Information private to the agent: disutility of effort.

Timing:

- (a) principal offers contract based on observable public information systems, including the output;
- (b) agent chooses action;
- (c) signals from the specified public information systems are observed;
- (d) agent gets paid on the basis of the signal;
- (e) a state of nature occurs;

- (f) output is observed;
- (g) principal keeps the residuum.

Special technological assumptions: Some of these assumptions are used in only some of the results; other results are obtained by relaxing them.

- (a) The joint probability distribution function on output, signals, and actions is twice-differentiable in effort, and the marginal effects on this distribution of the different components of effort are independent.
- (b) The principal's utility function U_p is twice differentiable, increasing, and concave.
- (c) The agent's utility function U_A is separable, with the function on the compensation scheme (or sharing rule as it is known) being increasing and concave, and the function on the effort being concave.

Results:

Result 2.1: There exists a marginal incentive informativeness condition which is essentially sufficient for marginal value given a signal information system Y . When information about the output is replaced by signals about the output and/or the agent's effort, marginal incentive informativeness is no longer a necessary condition for marginal value since an additional information system Z may be valuable as information about both the output and the effort.

Result 2.2: Information systems having no marginal insurance value but having marginal incentive informativeness may be used to improve risk sharing, as for example, when the signals which are perfectly correlated with output on the agent's effort are completely observable.

Result 2.3: Under the assumptions of result 2.2, when the output alone is observed, it must be used for both incentives and insurance. If the effort is observed as well, then a contract may consist of two parts: one part is based on the effort, and takes care of incentives; the other part is based on output, and so takes care of risk-sharing.

For example, consider auto insurance. The principal (the insurer) cannot observe the actions taken by the driver (such as care, caution and good driving habits) to avoid collisions. However, any positive signals of effort can be the basis of discounts on insurance premiums, as for example when the driver has proof of regular maintenance and safety check up for the vehicle or undergoes safe driving courses. Also factors such as age, marital status and expected usage are taken into account. The "output" in this case is the driving history, which can be used for risk-sharing; another indicator of risk which may be used is the locale of usage (country lanes or heavy city traffic). This example motivates result 2.4, a corollary to results 2.2 and 2.3.

Result 2.4: Information systems having no marginal incentive informativeness but having marginal insurance value may be used to offer improved incentives.

Result 2.5: If the uncertainty in the informative signal system is influenced by the choices of the principal and the agent, then such information systems may be used for control in decentralized decision-making.

5.3.3 Model 3

Holmstrom's model (Holmstrom, 1979) examines the role of imperfect information under two conditions: (i) when the compensation scheme is based on output alone, and (ii) when additional information is used. The assumptions about technology, information and timing are more or less standard, as in the earlier models. The model specifically uses the following:

- (a) In the first part of the model, almost all information is public; in the second part, asymmetry is brought in by assuming extra knowledge on the part of the agent.
- (b) output is a function of the agent's effort and state of nature: $q = q(e, \theta)$, and $\partial q / \partial e \geq 0$.
- (c) The agent's utility function is separable in compensation and effort, where $U_A(c)$ is defined on compensation, and $d(e)$ is the disutility defined on effort.
- (d) Disutility of effort $d(e)$ is increasing in effort.
- (e) The agent is risk averse, so that $U_A'' < 0$.
- (f) The principal is weakly risk neutral, so that $U_p'' \leq 0$.
- (g) Compensation is based on output alone.
- (h) Knowledge of the probability distribution on the state of nature θ is public.
- (i) Timing: The agent chooses effort before the state of nature is observed.

The problem:

$$(P) \max_{c \in C, e \in E} E[U_p(q - c(q))]$$

such that

$$E[U_A(c(q), e)] \geq \bar{U}, \quad (\text{IRC})$$

$$e \in \operatorname{argmax}_{e' \in E} E[U_A(c(q), e')]. \quad (\text{ICC})$$

To obtain a workable formulation, two further assumptions are made:

- (a) There exists a distribution induced on output and effort by the state of nature, denoted $F(q, e)$, where $q \equiv q(e, \theta)$. Since $\partial q / \partial e \geq 0$ by assumption, it implies $\partial F(q, e) / \partial e \leq 0$. For a given e , assume $\partial F(q, e) / \partial e < 0$ for some range of values q .
- (b) F has density function $f(q, e)$, where (denoting $f_e \equiv \partial f / \partial e$) f_e and f_{ee} are well defined for all (q, e) .

The ICC constraint in (P) is replaced by its first order condition using f , and the following formulation is obtained:

$$(P') \operatorname{Max}_{c \in C, e \in E} \int U_P(q - c(q)) f(q, e) dq$$

such that

$$\int [U_A(c(q)) - d(e)] f(q, e) dq \geq \bar{U}, \quad (\text{IRC}')$$

$$\int U_A(c(q)) f_e(q, e) dq = d'(e). \quad (\text{ICC}')$$

Results:

Result 3.1: Let λ and μ be the Lagrange multipliers for IRC' and ICC' in (P') respectively. Then, the optimal compensation schemes are characterized as follows:

$$\frac{U'_P(q - c(q))}{U'_A(c(q))} = \lambda + \mu \cdot \frac{f_e(q, e)}{f(q, e)} \quad a.e. [\underline{c}, \bar{c}],$$

where \underline{c} is the agent's wealth, and \bar{c} is the principal's wealth plus the output (these form the lower and upper bounds). If the equality in the above characterization does not hold, then $c(q) = \underline{c}$ or \bar{c} depending on the direction of inequality.

Result 3.2: Under the given assumptions and the characterization in result 3.1, $\mu > 0$; this is equivalent to saying that the principal prefers the agent increase his effort given a second-best compensation scheme as in the above result 3.1. The second-best solution is strictly inferior to a first-best solution.

Result 3.3: $|f_e|/f$ is interpreted as a benefit-cost ratio for deviation from optimal risk sharing. Result 3.1 states that such deviation must be proportional to this ratio taking individual risk aversion into account. From Result 3.2, incentives for increased effort are preferable to the principal. The following compensation scheme accomplishes this (where $c^F(q)$ denotes the first-best solution for a given λ):

$$\begin{aligned} c(q) &\geq c^F(q), \text{ if the marginal return on effort is positive to the agent;} \\ c(q) &< c^F(q), \text{ otherwise.} \end{aligned}$$

Result 3.4: Intuitively, the agent carries excess responsibility for the output. This is implied by result 3.3 and the assumptions on the induced distribution f .

A previous assumption is now modified as follows: Compensation c is a function of output and some other signal y which is public knowledge. Associated with this is a joint distribution $F(q, y, e)$ (as above), with $f(q, y, e)$ the corresponding density function.

Result 3.5: An extension of result 3.1 on the characterization of optimal compensation schemes is as follows:

$$\frac{U'_P(q - c(q,y))}{U'_A(c(q,y))} = \lambda + \mu \cdot \frac{f'_e(q,y,e)}{f(q,y,e)}$$

where λ and μ are as in result 3.1.

Result 3.6: Any informative signal, no matter how noisy it is, has a positive value if costlessly obtained and administered into the contract.

Note: This result is based on rigorous definitions of value and informativeness of signals (Holmstrom, 1979).

In the second part of this model, an assumption is made about additional knowledge of the state of nature revealed to the agent alone, denoted z . This introduces asymmetry into the model. The timing is as follows:

- (a) the principal offers a contract c based on the output and an observed signal y ;
- (b) the agent accepts the contract;
- (c) the agent observes a signal z about Θ ;
- (d) the agent chooses an effort level;
- (e) a state of nature occurs;
- (f) agent's effort and state of nature yield an output;
- (g) sharing of output takes place.

We can think of the signal y as information about the state of nature which both parties share and agree upon, and the signal z as special post-contract information about the state of nature received by the agent alone.

For example, a salesman's compensation may be some combination of percentage of orders and a fixed fee. If both the salesman and his manager agree that the economy is in a recession, the manager may offer a year-long contract which does not penalize the salesman for poor sales, but offers above subsistence level fixed fee to motivate loyalty to the firm on the part of the salesman, and a clause thrown in which transfers a larger share of output than normal to the agent (i.e. incentives for extra effort in a time of recession).

Now suppose the salesman, as he sets out on his rounds, discovers that the economy is in an upswing, and that his orders are being filled with little effort on his part. Then the agent may continue to exert little effort, realize high output, get a higher share of output in addition to a higher initial fixed fee as his compensation.

In the case of asymmetric information, the problem is formulated as follows:

$$(PA) \text{ Max}_{c(q,y) \in C, e(z) \in E} \int U_P(q - c(q,y))f(q,y|z,e(z))p(z)dqdydz$$

such that

$$\int U_A(c(q,y))f(q,y|z,e(z))p(z)dqdydz - \int d(e(z))p(z)dz \geq \bar{U}, \quad (IRC)$$

$$e(z) \in \text{argmax}_{e' \in E} \int U_A(c(q,y))f(q,y|z,e')dqdy - d(e') \quad \forall z \quad (ICC)$$

where $p(z)$ is the marginal density of z , $d(e(z))$ is the disutility of effort $e(z)$.

Let λ and $\mu(z)p(z)$ be the Lagrange multipliers for (IRC) and (ICC) in (PA) respectively.

Result 3.7: The extension of result 3.1 on the characterization of optimal compensation schemes to the problem (PA) is:

$$\frac{U'_P(q - c(q,y))}{U'_A(c(q,y))} = \lambda + \frac{\int \mu(z) f_e(q,y|z,e(z)) p(z) dz}{\int f(q,y|z,e(z)) p(z) dz}$$

The interpretation of result 3.7 is similar to that of result 3.1. Analogous to result 3.2, $\mu(z) \neq 0$, and $\mu(z) < 0$ for some z and $\mu(z) > 0$ for other z , which implies, as in result 3.2, that result 3.7 characterizes solutions which are second-best.

5.3.4 Model 4: Communication under Asymmetry

This model (Christensen, 1981) attempts an analysis similar to model 3, and includes communication structures in the agency. The special assumptions are as follows:

- (a) There is a set of messages M that the agent uses to communicate with the principal; compensation is based on the output and the message picked by the agent; hence, the message is public knowledge.
- (b) There is a set of signals about the environment; the agent chooses his effort level based on the signal he observes; the agent also selects his compensation scheme at this time by selecting an appropriate message to communicate to the principal; selection of the message is based on the effort.

- (c) Uncertainty is with respect to the signals observed by the agent; the distribution characterizing this uncertainty is public knowledge; the joint density is defined on output and signal conditioned on the effort:

$$f(q, \xi | e) = f(q | \xi, e) \cdot f(\xi).$$

- (d) Both parties are Savage(1954)-rational.
- (e) The principal's utility of wealth is U_p , with weak risk-aversion; in particular,

$$U_p' > 0 \quad \text{and} \quad U_p'' \leq 0.$$

- (f) The agent's utility of wealth is separable into U_A defined on compensation and disutility of effort. The agent has positive marginal utility for money, and he is strictly risk-averse; i.e. $U_A' > 0$, $U_A'' < 0$, and $d' > 0$.

Timing:

- (a) The principal and the agent determine the set compensation schemes, based on the output and the message sent to the principal by the agent; the principal is committed to this set of compensation schemes;
- (b) the agent accepts the compensation scheme if it satisfies his reservation welfare;
- (c) the agent observes a signal ξ ;
- (d) the agent picks an effort level based on ξ ;
- (e) the agent sends a message m to the principal; this causes a compensation scheme from the contracted set to be chosen;
- (f) output occurs;

(g) sharing of output takes place.

Note that in the timing, (d) and (e) could be interchanged in this model without affecting anything.

The following is the principal's problem:

(P) Find $(c^*(q,m), e^*(\xi, m), m^*(\xi))$ such that $c^* \in C$, $e^* \in E$, and

$m^* \in M$ solves:

$$\text{Max}_{c(q,m), e(\xi), m(\xi)} E[U_P(q - c(q,m))]$$

such that

$$E[U_A(c(q,m)) - d(e)] \geq \bar{U}, \quad (\text{IRC})$$

$$e(\xi) \in \text{argmax}_{e' \in E} E[U_A(c(q, m(\xi))) - d(e')] \quad (\text{self-selection of action}),$$

$$m(\xi) \in \text{argmax}_{m' \in M} E[U_A(c(q, m')) - d(e(\xi, m')) | \xi] \quad (\text{self-selection of message}),$$

where $e(\xi, m)$ is the optimal act given that ξ is observed and m is reported.

The following assumptions are used for analyzing the problem in the above formulation:

- (a) $U_P(\cdot)$ and $U_A(\cdot) - d(\cdot)$ are concave and twice continuously differentiable in all arguments.
- (b) Compensation functions are piecewise continuous and differentiable a.e. (ξ) .
- (c) The density function f is twice differentiable a.e.
- (d) Regularity conditions enable differentiation under the integral sign.
- (e) Existence of an optimal solution is assumed.

Result:

Result 4.1: The following is a characterization of optimal functions:

$$\frac{U'_P(q - c^*(q, \xi))}{U'_A(c^*(q, \xi))} = \lambda - \rho'(\xi) + \frac{(\mu(\xi) - \rho(\xi)) \frac{\partial}{\partial \xi} e^*(\xi) f_e(q, \xi | e^*(\xi))}{f(q, \xi | e^*(\xi))} - \frac{\rho(\xi) f_\xi(q, \xi | e^*(\xi))}{f(q, \xi | e^*(\xi))}$$

where λ , $\mu(\xi)$, and $\rho(\xi)$ are Lagrange multipliers for the three constraints in (P) respectively.

5.3.5 Model G: Some General Results

Result G.1 (Wilson, 1968). Suppose that both the principal and the agent are risk averse having linear risk tolerance functions with the same slope, and the disutility of the agent's effort is constant. Then the optimal sharing rule is a non-constant function of the output.

Result G.2. In addition to the assumptions of result G.1, also suppose that the agent's effort has negative marginal utility. Let $c_1(q)$ be a sharing rule (or compensation scheme) which is linear in the output q , and let $c_2(q) = k$ be a constant sharing rule. Then, c_1 dominates c_2 .

The two results above deal with conditions when observation of the output is useful. Suppose Y is a public information system that conveys information about the output. So, compensation schemes can be based on Y alone. The value of Y , denoted $W(Y)$ (following model 1), is defined as: $W(Y) \equiv \max_{c \in C} EU_P[q - c(y)]$, subject to IRC

and ICC. Let Y^0 denote a non-informative signal. Then, the two results yield a ranking of informativeness: $W(Y) > W(Y^0)$. When Q is an information system denoting perfect observability of the output q , and the timing of the agency relationship is as in model 1 (i.e. payment is made to the agent after observing the output), then $W(Q) > W(Y)$ as well.

CHAPTER 6

METHODOLOGICAL ANALYSIS

The solution to the principal-agent problem is influenced by the way the model itself is setup in the literature. Highly specialized assumptions, which are necessary in order to use the optimization technique, contribute a certain amount of bias. As an analogy, one may note that a linear regression model assumes implicit bias by seeking solutions only among linear relationships between the variables; a correlation coefficient of zero therefore implies only that the variables are not linearly correlated, not that they are not correlated. Examples of such specialized assumptions abound in the literature, a small but typical sample of which are detailed in the models presented in Chapter 5.

The consequences of using the optimization methodology are primarily of two. Firstly, much of the pertinent information that is available to the principal, the agent and the researcher must be ignored, since this information deals with variables which are not easily quantifiable, or which can only be ranked nominally, such as those that deal with behavioral and motivational characteristics of the agent and the prior beliefs of the agent and principal (regarding the task at hand, the environment, and other exogenous variables). Most of this knowledge takes the form of rules linking antecedents and consequences, and which have associated certainty factors.

Secondly, a certain amount of bias is introduced into the model by requiring that the functions involved in the constraints satisfy some properties, such as differentiability, monotone likelihood ratio, and so on. It must be noted that many of these properties are reasonable and meaningful from the standpoint of accepted economic theory. However, standard economic theory itself relies heavily on concepts such as utility and risk aversion in order to explain the behavior of economic agents. Such assumptions have been criticized on the grounds that individuals violate them; for example, it is known that individuals sometimes violate properties of the Neumann-Morgenstern utility functions. Decision theory addressing economic problems also uses concepts such as utility, risk, loss, and regret, and relies on classical statistical inference procedures. However, real life individuals are rarely consistent in their inference, lacking in statistical sophistication, and unreliable on probability calculations. Several references to support this view are cited in Chapter 2. If the term "rational man" as used in economic theory means that individuals act as if they were sophisticated and infallible (in terms of method and not merely content), then economic analysis might very well yield erroneous solutions.

Consider, as an example, the treatment of compensation schemes in the literature. They are assumed to be quite simple, either being linear in the output, or involving a fixed element called the rent. (See chapter 5 for details). In practice, compensation schemes are fairly comprehensive and involved. They cover as many contingencies as possible, provide for a variety of payment and reward criteria, specify grievance procedures, termination, promotion, varieties of fringe benefits, support services, access to company resources, and so on.

The set of all compensation schemes is in fact a set of knowledge bases consisting of the following components (B.R. Ellig, 1982):

- (1) Compensation policies/strategies of the principal;
- (2) Knowledge of the structure of the compensation plans, which means specific rules concerning short-term incentives linked to partial realization of expected output, long-term incentives linked to full realization of expected output, bonus plans linked to realizing more than the expected output, disutilities linked to underachievement, and rules specifying injunctions to the agent to restrain from activities that may result in disutilities to the principal (if any).

There are various elements in a compensation scheme, which can be classified as financial and non-financial:

Financial elements of compensation

1. Base Pay (periodic).
2. Commission or Share of Output.
3. Bonus (annual or on special occasions).
4. Long Term Income (lump sum payments at termination).
5. Benefits (insurance, etc.).
6. Stock Participation.
7. Non-taxable or tax-sheltered values.

Nonfinancial elements of compensation

1. Company Environment.
2. Work Environment.

3. Items which are designed to improve productivity of agent.
4. Status or Prestige.
5. Elements of agent's disutility assumed by the firm.

As another example, note that some of the important factors not considered in the traditional treatment of the principal-agent problem are connected to the characteristics of the agent. In a real-world situation, the principal has a great deal of behavioral knowledge which he acquires from acting in a social context. In dealing with the problems associated with the agency contract, he takes into account factors of the agent such as the following:

- * General social skills, which are also known as social interaction skills, networking skills, or people skills.
- * Office and managerial skills.
- * Past experience or reputation.
- * Motivation or enthusiasm.
- * General behavioral aspects (personal habits).
- * Physical qualities deemed essential or useful to the task.
- * Language/communication skills.

In the light of these shortcomings of the traditional methodology, it is desirable to see how they make their decisions in reality. It may be more fruitful to think of people making decisions based on some underlying probabilistic knowledge bases. These knowledge bases would capture all the rules of behavior and decision-making, such as

the choice of effort levels by the agent. This would also enable one to bypass the use of utility and risk aversion as artificial explanatory variables.

In order to see how behavioral and motivational factors may be integrated into the new approach, it is necessary to review briefly some models of motivation and behavioral theory. This is done in Chapter 7.

CHAPTER 7 MOTIVATION THEORY

There are many models of motivation. One is drive theory (W.B. Cannon, 1939; C.L. Hull, 1943). The main assumption in drive theory is that decisions concerning present behavior are based in large part on the consequences, or rewards, of past behavior. Where past actions led to positive consequences, individuals would tend to repeat such actions; where past actions led to negative consequences or punishment, individuals would tend to avoid repeating them. C.L. Hull (1943) defines "drive" as an energizing influence which determined the intensity of behavior, and which theoretically increased along with the level of deprivation. "Habit" is defined as the strength of relationship between past stimulus and response (S-R). The strength of this relationship depends not only upon the closeness of the S-R event to reinforcement but also upon the magnitude and number of such reinforcements. Hence effort, or motivational force, is a multiplicative function of magnitude and number of reinforcements.

In the context of the principal-agent model, drive theory would explain the agent's effort as arising from some past experience of deprivation (need of money) and from the strength of feeling that effort leads to reward. So, the drive model of motivation defines effort as follows:

Effort = Drive * Habit

$$= f_d(\text{past deprivation}) * f_h(\Sigma | S-R |)$$

where f_d is some "function" denoting drive as dependent on past deprivation,

f_h is some "function" denoting habit as dependent on the sum of the strengths of a number of instances of S-R reinforcements, and

$| S-R |$ is the magnitude of an S-R reinforcement.

Drive theory, in its simplest form, states that individuals have basic biological drives (eg., hunger and thirst) that must be satisfied. As these drives increase in strength, there is an accompanying increase in tension. Tension is aversive to the organism, and anything reducing that tension is viewed positively. The process of performing action that achieves this is termed **learning**. All higher human motives are deemed to be derivatives of this learning.

Another view is given in Instrumentality Theory which rejects the drive model (L.W. Porter & E.E. Lawler, 1968), and emphasizes the anticipation of future events. This emphasis provides a cognitive element ignored in most of the drive models. The reasons for preferring instrumentality theory over other theories may be summarized as follows:

- (1) The terminology and concepts of instrumentality theory are more applicable to the problems of human motivation; the emphasis on rationality and cognition is appropriate for describing the behavior of managers.
- (2) Instrumentality theory greatly facilitates the incorporation of motives such as status, achievement, and power into a theory of attitudes and performance.

Figure 1 shows the Porter & Lawler model of the instrumentality theory of motivation. The model parts are described below.

Value of reward describes the attractiveness of various outcomes to the individual. The instrumentality model agrees with the drive model that rewards acquire attractiveness as a function of their ability to satisfy the individual.

Perceived effort-reward probability refers to the subjective estimate of the individual that increased effort will lead to the acquisition of some valued reward. This consists of two estimates: the first is the probability that improved performance will lead to the value reward, and the second is the probability that effort will lead to improved performance. These two probabilities have a multiplicative relationship. Instrumentality model makes a distinction between **Effort** and **Performance**: effort is a measure of how hard an individual works, while performance is a measure of how effective is his effort.

Abilities and traits are included as a source of variation in this model, while other models implicitly assume some fixed levels of abilities and traits. Abilities and traits refer to relatively stable characteristics of the individual such as intelligence, personality characteristics, and psychomotor skills, which are considered as boundary conditions or limitations on performance.

Role Perception denotes an individual's definition of successful performance in work. An appropriate definition of success is essential in determining whether or not effort is transformed into good performance, and also in perceiving equity in reward.

Distinction is made between intrinsic and extrinsic rewards. Intrinsic rewards are rewards that satisfy higher-order Maslow needs (A.H. Maslow, 1943; A.H. Maslow,

1954) and are administered by the individual to himself rather than by some external agent. Extrinsic rewards are rewards administered by an external party such as the principal.

Perceived equitable rewards describes the level of reward that an individual feels is appropriate. The appropriateness of the reward is linked to role perceptions and perception of performance.

Satisfaction is referred to as a "derivative variable". It is derived by the individual (here, the agent) by comparing actual reward to perceived equitable reward. Satisfaction may therefore be defined as the correspondence or correlation between actual reward and perceived equitable reward.

Research in instrumentality theory is detailed in (Campbell and Pritchard, 1976; Mitchell, 1974). Most of the tests of both their initial model and later versions have yielded similar results: effort is predicted more accurately than performance. This makes sense logically. Individuals have effort under their control but not always performance. The environment (exogenous or random risk) plays a major role in determining if and how effort yields levels of performance (Steers and Porter, 1983).

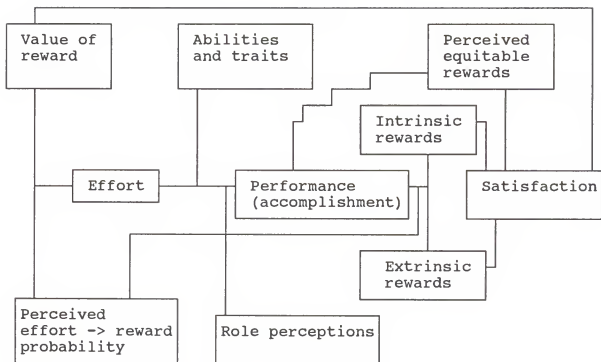


FIGURE 1: THE PORTER AND LAWLER MODEL OF INSTRUMENTALITY THEORY

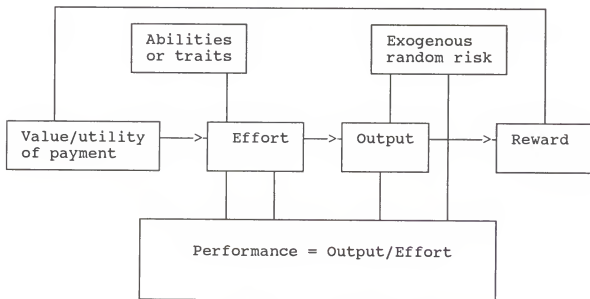


FIGURE 2: MODIFIED PORTER AND LAWLER MODEL

CHAPTER 8

RESEARCH FRAMEWORK

The object of the research is to develop and demonstrate an alternative methodology for studying agency problems. To this end, we study several agency models from a common framework described below. There are two types of issues associated with the studies. One deals with the issues of modeling the agency problem itself. The other deals with the issues of the method, in this case, knowledge bases, genetic learning operators, and the operators of specialization and generalization.

The common framework for the agency problems has these elements:

1. The use of rule bases to model the information and expertise possessed by the principal and the agent.
2. The use of probability distributions to model the uncertain nature of some of the information.
3. Consideration of a number of elements of compensation.
4. Offering compensation to an agent based on the agent's characteristics.

The common framework for the methodology for studying agency issues has these elements:

1. Simulation of the agency interactions over a period of time.

2. Use of learning mechanisms to capture the dynamics of agency interaction.

A number of preliminary studies were conducted in order to define and fine-tune the two frameworks. The initial studies sought to understand the behavior of optimal compensation schemes in a dynamic environment. These initial studies supported the idea that learning by way of the genetic algorithm paradigm leads to quick convergence to a relatively stable solution. Of course, genetic algorithms may find multiple stable solutions. Further, the preliminary studies led to fixing of the genetic parameters, since it was noticed that variations in these parameters did not contribute anything of interest. For example, increasing mutation probability delayed convergence of the solutions, and beyond 0.5 led to a chaotic situation. Similarly, varying the mating probability had an effect on the speed with which the solutions were found. The nature of the solutions were not affected. The genetic parameters were therefore fixed as follows:

- * Crossover mechanism - uniform one-point crossover;
- * Mating probability - 0.6;
- * Mutation probability - ranging from 0.01 to 0.001 (for different models);
- * Discard worst rule and copy the best rule.

The use of generalization and specialization operators for learning in later models will be described subsequently. Below, we give an overview of the various models studied. The details follow in later sections.

Models 1 and 2 were preliminary studies conducted to explore the new framework for attacking agency problems. The goal of these models, as well as Model 3, was to demonstrate the feasibility of addressing issues in agency which the traditional theory

ignored (see, for example, Chapter 6 for a methodological analysis). Models 1 and 2 led to the choice of genetic parameters (as described above), and finalizing the agency interaction mechanism (namely, timing and information), including the Porter-Lawler model of human behavior and motivation. While both Models 1 and 2 are more realistic than the traditional models, they still do not capture the entire realism of an agency. The later models capture increasing amounts of realism.

Model 3 is the first formal study. The goal of this study is to develop a model which provides a counter-example to the traditional theory which considers fixed pay, share of output, and exogenous risk to be important agency variables, and ignores the role of the agent's behavioral and motivational characteristics in selecting his compensation scheme. This study tries to answer the following questions: Is there a non-trivial and formal agency scenario where the lack of dependence of the compensation scheme on the agent's characteristics leads to a sub-optimal solution (as compared to the standard theory)? Is there a scenario wherein consideration of other elements of compensation lead to better results for both the principal and the agent? Is there a scenario where, from a principal's perspective, exogenous risk (which can only be observed ex-post) plays a lesser role than other agency variables? How does certainty of information affect the nature of the solutions? What measures may be used to characterize good solutions, or identify important variables? The last question is non-trivial, because all the variables used in these studies are discrete nominal valued, and hence are not amenable to any formal measure theory. This study involves five experiments (which differ in the information available to the principal), and the use of

factor analysis to study the principal's knowledge base at the end of the simulation in order to characterize good compensation schemes and identify important variables.

Models 1, 2 and 3 involve only a single agent. Models 4 and beyond capture more realism in agency relationship. They are multi-agent, multi-period, dynamic (agents are hired and fired all the time) models. Moreover, they closely follow one traditional agency theory - the LEN model of Spemann (see Chapter 5 for details). Models 4 and 5 study the LEN model, while including only two elements of compensation as in the original LEN model, and retaining the behavioral characteristics of the agents. Model 4 studies the agency under a non-discriminatory firing policy of the principal, while Model 5 studies exactly the same agency but with the principal employing a discriminatory firing policy for the agents. Similarly, Models 6 and 7 are non-discriminatory and discriminatory respectively. However, Models 6 and 7 employ compensation variables not included in the original LEN model. These models study the following issues:

- * the nature of good compensation schemes under a demanding agency environment;
- * the correlation between various variables of the agency;
- * the correlation between the variables of the agency and the control variables of the experiments;
- * the effect of discriminatory firing practices by the principal
- * the effect of complex compensation schemes.

Chapter 9 describes Model 3 in detail. Chapter 10 introduces Models 4 through 7, and describes each in detail. The conclusions are given in Chapter 11, and directions for future research are covered in Chapter 12.

CHAPTER 9

MODEL 3

9.1 Introduction

In Model 3, utility functions are replaced by knowledge bases, machine learning replaces estimation and inference replaces optimization. In so doing, complex contractual structures and behavioral and motivational considerations can be directly incorporated into the model.

In Section 2 we describe a series of experiments used to illustrate our approach. These experiments study a realistic situation. Section 3 covers the methodology and details of the experiments. Section 4 tabulates the results of our experiments, while Section 5 describes and discusses the results.

Initially, the principal's knowledge base reflects her current state of knowledge about the agent (if any). The agent's knowledge base reflects the way he will produce under a contract. This knowledge base incorporates motivational and behavioral characteristics. It includes his perception of exogenous risk, social skills, experience, etc. The details are provided in Section 3.

The principal will refine her knowledge base through a learning mechanism. Using the current knowledge base, the principal will use inference to determine a

contract. This contract is used by the agent. The resulting output and welfare are used by the principal to construct a "better" knowledge base through a learning procedure.

In the following we incorporate specific models and components to achieve an implementation of our new principal-agent model. We link behavioral factors by the model of Porter & Lawler (1968), which also incorporates the calculation of satisfaction and subsequent effort levels by the agent. The Porter & Lawler model derives from the instrumentality theory of motivation, which emphasizes the anticipation of future events, unlike most models of motivation based on drive theory. The key ideas of the Porter & Lawler model are the recognition of the appropriateness of rationality and cognition as descriptive of the behavior of managers, and the incorporation of motives such as status, achievement, and power as factors that play a role in attitudes and performance.

Effort is determined by the utility or value of compensation and the perceived probability of effort leading to reward. Performance, determined by the effort level, abilities of the agent, and role perceptions, leads to intrinsic and extrinsic rewards, which in turn influence the satisfaction derived by the agent. A comparison of performance and the satisfaction derived from it influences the perception of equity of reward, and reinforces or weakens satisfaction. Performance also plays a role in the revision of the probability of effort leading to adequate reward.

The principal and agent knowledge-bases in our model consist of rules. Each rule has a set of antecedent variables and a set of consequent variables. The antecedent variables are the agent's behavioral characteristics and the exogenous risk, while the consequent variables are the variables denoting the elements of compensation. The

knowledge-base therefore consists of rules that specify selection of compensation plans based on the agent's characteristics and the exogenous risk. A formal description follows.

Let N denote a nominal scale. N^k denotes the k -fold nominal product. Let $C \subseteq N^k$ denote the set of all compensation plans, $\{c_1, \dots, c_n\}$. Let $B \subseteq N^k$ denote the set of all the behavioral profiles of the agent, $\{b_1, \dots, b_m\}$. Each compensation plan c_i is a finite-dimensional vector $c_i \equiv (c_i^{(1)}, \dots, c_i^{(p)})$, where each element of the c_i vector denotes an element of the compensation plan, such as fixed pay, commission, or bonus. Each element of B is also a finite-dimensional vector, $b_j \equiv (b_j^{(1)}, \dots, b_j^{(q)})$, where each element of the b_j vector denotes a behavioral characteristic that the agent will be evaluated on by the principal, such as experience, motivation, or communication skill. The elements c_i^0 and b_j^0 are detailed in Section 5.

Let G be the set of mappings from the set of behavioral profiles B to the set of compensation plans, C . Two or more compensation plans could conceivably be associated with one particular behavioral profile b_j of an agent. A particular mapping g in G specifies a knowledge base $K \subseteq B * C$.

Let $S: K * \Theta * E \rightarrow \mathbb{R}$ denote the total satisfaction function of the agency, where E denotes the effort level of the agent, and Θ represents exogenous risk. Let S_A denote the satisfaction of the agent, and S_p , the satisfaction of the principal (defined on the same domain). S_A and S_p are both "functions" of other variables such as compensation plans, output, agent's effort, agent's private information, and so on.

$$S_A \equiv S_A(\text{Output}, C),$$

$$S_p = S_p(V, \text{Effort}, C), \text{ and}$$

$$S = S(S_A, S_p),$$

where V is the agent's private information about the principal and her company.

Thus, $S(b_i, c_i)$ denotes the total satisfaction derived when the agent has the behavioral profile b_i and the principal offers compensation plan c_i . Define fitness to be the total satisfaction $S(b_i, c_i)$ normalized with respect to the whole knowledge base K . Let $F(g)$ denote the average fitness of a mapping $g \in G$ which specifies a knowledge base $K_g \subseteq B * C$.

$$F(g) = \frac{1}{n} \sum_{i=1}^n S(b_i, c_i), \quad b_i \in B, \quad c_i \in C.$$

The objective function of the principal-agent problem in our formulation is:

$$\begin{aligned} \max_{g \in G} \quad & E[F(g)] = E\left[\frac{1}{n} \sum_{i=1}^n S(b_i, c_i)\right] \\ = \quad & E\left[\frac{1}{n} \sum_{i=1}^n S(B, C, \Theta, V)\right] \\ = \quad & E[\overline{S}(E, C, \Theta, V)]. \end{aligned}$$

Our formulation of the principal-agent problem may be stated formally as:

$$\max_{g \in G} \quad E[F(g)] = E[\overline{S}(E, C, \Theta)]$$

such that

$$E \in \operatorname{argmax} \overline{S}_A(B, C, \Theta, V),$$

where

$C \in \text{range}(g)$, and

V is the agent's private information.

The two constraints of the original problem (Individual Rationality and Incentive Compatibility) are subsumed in the calculation of F . The agent, for example, selects his effort level so as to increase his satisfaction or welfare based on his behavioral characteristics and the compensation plan offered by the principal. It is not necessary to check for IRC explicitly. Our model ensures that the agent, when presented with a compensation plan that does not satisfy his IRC, picks an effort level that yields extremely low total satisfaction. The dynamic learning process (described below) discards all such compensation plans. In order to formalize the constraints in our new model, it is necessary to introduce details of the functions, knowledge bases, representation scheme for the knowledge bases, and the inference strategy. This is done in Section 9.3.

9.2 An Implementation and Study

To both illustrate our method and to study the results of our approach, a series of experiments were conducted. All the simulation experiments start with the same initial set of rules for the principal, with the variables denoting agent characteristics acting as the antecedents and the variables denoting elements of compensation acting as consequents. This initial knowledge base of 500 rules is generated randomly, which ensures that no initial bias is introduced into the model.

The agent's knowledge base is varied from experiment to experiment to reflect different behavioral characteristics, abilities, and perceptions. The experiments differ with respect to each other in the probability distributions of the variables representing the agent's characteristics and the agent's personal information about the principal.

An experiment consists of 10 runs of a sequence of 200 learning cycles including the following steps:

1. Using her current knowledge base, the principal infers a compensation plan.
2. The agent performs under this compensation plan and an output is realized.
3. A satisfaction level is computed which reflects the total welfare of the principal and the agent.
4. The principal notes the results of the compensation plan and revises her knowledge base using a genetic algorithm learning method.

The following hypotheses are considered:

Hypothesis 1: Behavioral characteristics and complex compensation plans play a significant role in determining good compensation rules.

Hypothesis 2: In the presence of complete certainty regarding behavioral characteristics, the most important variables that explain variation in good compensation rules are the same as those considered in the traditional principal-agent models.

Hypothesis 3: Extra information about behavioral characteristics yields better compensation rules. Specifically, any information is better than having non-informative priors.

The experiments are designed to study the compensation rules which achieve close-to-optimal satisfaction for the principal and the agent under different informational assumptions. Each experiment pertains to a different agent having specific behavioral characteristics and perception of the principal or the company.

Nine characteristics of the agent. They are: experience, education, age, general social skills, office and managerial skills, motivation, physical qualities deemed essential to the task, language and communication skills, and miscellaneous personal characteristics.

The elements of compensation that are taken into account are: basic pay, share of output or commission, bonus payments, long term payments, benefits, and stock participation.

In the calculation of satisfaction (total welfare of the principal and the agent), we also take into account variables that denote the agent's perception or assessment of the principal or her company. These variables may be called the agent's "personal" variables, since the principal has no information about them. The agent's personal variables we consider are: company environment, work environment, status, his own traits and abilities, and his perceived probability of effort leading to reward.

Characterization of the agent in each of the five experiments is given below: Agent #1 (involved in Experiment 1) is moderately experienced, has completed high school, is above 55 years of age. His general social skills are average, but his office and managerial skills are quite good. He has slightly above average motivation and enthusiasm for the job, and he is more or less physically fit, but the principal is not very

sure about the agent's health. He has good communication skills, while his other miscellaneous personal characteristics leave something to be desired. He has a rather pessimistic outlook about exogenous factors (economy and business conditions), but he is not too sure of his pessimistic estimate.

His assessment of company and work environment is moderately favorable, while he considers the company's corporate image to be rather high. His perception of his own abilities and traits is that they are just better than average, but he feels that he is not consistent (sometimes he does far better, sometimes far worse). He is pessimistic about effort leading to reward, perhaps because his uninspiring characteristics and lack of good education led to slow reward and promotion in the past.

Agent #2 in Experiment 2 has the same assessment of personal variables as Agent #1. However, his characteristics are more modest. He has very little experience, no high school education, and is much below average in all other respects. He is as pessimistic and as unsure about the exogenous variables as Agent #1.

Agent #3 in Experiment 3 is a college graduate in his late 20's to early 30's. He has little experience but is very highly motivated, possesses good communication skills, and is good in all the other characteristics. His assessment of the exogenous environment is optimistic. Moreover, he believes the principal's work and company environment is very good, and is generally sure of his superior abilities. He believes effort will almost always be rewarded appropriately.

Nothing is known about Agent #4 in Experiment 4, while everything known about Agent #5 in Experiment 5 is known with certainty. Agent #5 is in the same age bracket as Agent #3, while he has more experience. His office and managerial skills, motivation and enthusiasm are of the highest. He is physically very fit, and has very good communication skills. He perceives the principal's company and work environment to be the best in the market, and he is very certain of his superior talents. He firmly believes that effort is always rewarded.

The nominal scales are used throughout the experiments are given below.

For the variables CE, WE, SI, AT, GSS, OMS, P, L, OPC:

1: very bad, 2: bad, 3: average, 4: good, 5: excellent.

For PPER and M (Motivation):

1: very low, 2: low, 3: average, 4: high, 5: very high.

For X (Experience):

1: none, 2: < 1 year, 3: between 1 and 5 years,

4: between 5 and 10 years, 5: more than 10 years.

For D (Education):

1: below high school, 2: high school, 3: undergraduate,

4: graduate, 5: graduate (specialization/2 or more degrees).

For A (Age):

1: < 18 years, 2: between 18 and 25 years, 3: between 25 and 35 years, 4:

between 35 and 50 years, 5: above 50 years.

For RISK:

1: very high, 2: high, 3: average, 4: low, 5: very low.

Table 9.1 provides details that capture the above characterizations.

The information on each variable in Table 9.1 is specified as a discrete probability distribution. Table 9.1 lists the means and standard deviations of the variables associated with the agent's characteristics and the agent's personal variables. In experiment 4, the situation is non-informative. All the variables have discrete uniform distribution, with mean 3.00 and standard deviation $\sqrt{2} \approx 1.414$. Experiment 5 is provided complete and perfect information, and so the standard deviation is 0.00.

9.3 Details of Experiments

In this Section we discuss rule representation, inference method, calculation of satisfaction, details of the genetic learning algorithm, and statistics captured for analysis.

9.3.1 Rule Representation

A rule has the following format:

IF <antecedent> THEN <consequent> .

The antecedent values in the "IF" part of a rule are conditions that occur or are satisfied, and the consequent variables are assigned values from the "THEN" part of the rule correspondingly. The antecedent and consequent of a rule are conjunctions of several variables. Let b_i be the i -th antecedent (denoting a behavioral variable), and c_j denote the j -th consequent (denoting a compensation variable). The antecedent of a rule is then given by $\bigwedge_{i=1,\dots,m}$ and the consequent by $\bigwedge_{j=1,\dots,n}$, where \bigwedge denotes conjunction. Hence,

for a rule to be activated in our experiments, all the specified antecedent conditions must be fulfilled, and the result of the activation of the rule is to yield a compensation plan having all the specified elements. Hence, a compensation plan is dependent on the specific characteristics of the agent and also on the exact realization of exogenous risk. The effectiveness of each compensation plan is therefore dependent on how well it takes into account the characteristics of the agent. It is not necessary for each rule in the knowledge base to have all the m antecedents and all the n consequents specified. However, we adopt a uniform representation for the knowledge base where all the rules have full specification of all the antecedent and consequent variables.

All the variables are positionally fixed, which facilitates pattern-matching during inference (described in Sec. 5.2 below). The antecedent variables dealing with the agent's characteristics (including exogenous risk) are listed in order below, with the variable names in parentheses (b_4 is not a behavioral variable: it represents Θ , the exogenous risk):

- (1) Experience (X),
- (2) Education (D),
- (3) Age (A),
- (4) Exogenous Risk (RISK),
- (5) General Social Skills (GSS),
- (6) Office and Managerial Skills (OMS),
- (7) Motivation (M),
- (8) Physical Qualities deemed essential to the task (PQ),

- (9) Language and Communication Skills (L), and
- (10) Miscellaneous Personal Characteristics (OPC).

The consequent variables that denote the elements of compensation plans are listed in order below, with the variable names in parentheses:

- (1) Basic Pay (BP),
- (2) Share or Commission of Output (S),
- (3) Bonus Payments (BO),
- (4) Long Term Payments (TP),
- (5) Benefits (B), and
- (6) Stock Participation (SP).

We assume that each of the 10 variables representing the agent's characteristics (including exogenous risk) and the 6 variables that represent the elements of compensation has 5 possible values. This is a convenient number of values for nominal variables and represents one of the Likert scales.

In effect, every rule is represented as an ordered sequence of 16 integer numbers of 1 through 5. The first ten numbers are understood to be the antecedents, and the next six the consequents. The nominal scale linked to the consequent variables is as follows:

1: minimum; 2: low; 3: average; 4: high; 5: very high

For example, consider the following rule:

IF <2,3,1,4,5,2,3,1,4,3> THEN <3,2,4,3,2,2>

This rule means:

IF

Experience is less than one year, AND

Education is undergraduate, AND

Age is below 18 years, AND

Exogenous RISK is low (favorable business climate), AND

General Social Skills are excellent, AND

Office and Managerial Skills are bad (no skills at all), AND

Motivation is average, AND

Physical Qualities are very bad (frail health), AND

Communication Skills are good, AND

Other Characteristics are good

THEN

Basic Pay is average, AND

Commission is low, AND

Bonus payments are high, AND

Long term payments are average, AND

Benefits are low, AND

Stock Participation is low.

The total number of possible rules for the principal is $5^{16} = 152,587,890,625$. The goal of each trial is to pick a small number, say 500 ($= 3.2768 * 10^{-7} \%$) of rules from among these 5^{16} rules so that the final rules have very high satisfaction associated with them.

9.3.2 Inference Method

The key heuristics that motivate the inference process are:

- (1) compensation plans are conditional on the characteristics of the agent and the assessment of exogenous risk;
- (2) compensation plans which are close to optimal, rather than optimal, are sought.

We assume that the agent and the principal both have the same information on the exogenous risk. At each learning episode in an experiment, the values in the rules are changed by means of applying genetic operators (see Chapter 3 for details). The learning algorithm ensures that rules having "robust" combinations of compensation plans survive and refine over learning episodes. Such compensation plans are then identified as most effective for that particular agent.

The "functional" relationship of the different variables in the inference scheme is as follows (the subscript t denotes the learning episode or time):

$$\text{Effort}_t = g(f_1(C_t), f_2(V_t), \text{PPER}_t, \text{IR}_{t-1}, \text{PEPR}_t),$$

where C_t is the compensation offered by the principal in time or learning episode t ,

where PPER denotes perceived probability of effort leading to reward,

PEPR denotes perceived equity of past reward,

PERR is perceived equity of current reward,

$V_t = (CE_t, WE_t, ST_t, AT_t)$, the agent's private information in time t ,

g is a fixed real-valued effort selection mapping,

f_1 is a fixed real-valued mapping of compensation, and

f_2 is a fixed real-valued mapping of the agent's private information; similarly, functions f_3 through f_{10} , and h_1 through h_3 are fixed real-valued mapping defined on the appropriate domains;

$$\text{Output}_t = f_3(\text{Effort}_t, \text{RISK}_t);$$

$$\text{PERF}_t = f_4(\text{Output}_t, \text{Effort}_t);$$

$$\text{IR}_t = f_5(\text{CE}_t, \text{WE}_t, \text{ST}_t);$$

$$\text{PERR}_t = f_6(\text{PERF}_t, h_1(C_t));$$

$$\text{Disutility}_t = f_7(\text{Effort}_t);$$

$$\text{PEPR}_t = \text{PERR}_{t-1};$$

$$S_{At} = f_8(\text{PERF}_t, \text{IR}_t, h_2(C_t), \text{PERR}_t, \text{Effort}_t, \text{RISK}_t, \text{Disutility}_t);$$

$$S_{Pt} = f_9(\text{Output}_t, h_3(C_t, \text{Output}_t)); \text{ and}$$

$$S_t = f_{10}(S_{At}, S_{Pt}).$$

The functions g , f_1 through f_{10} , and h_1 through h_3 used in the inference scheme to select effort levels, infer intrinsic reward, disutility, satisfactions, etc., are given in Section 5.3 below.

9.3.3 Calculation of Satisfaction

At each learning episode, the following steps are carried out to compute the satisfaction of the principal and the agent:

- (0) the principal infers a compensation plan;
- (1) the agent selects an effort level based on the compensation plan, his perception of the principal, and other variables from the Porter & Lawler model;

- (2) an act of nature is generated randomly according to the specified distribution of exogenous risk;
- (3) output is a function of effort and the act of nature;
- (4) performance is a function of output and effort;
- (5) the agent's intrinsic reward is calculated;
- (6) the agent's perceived equity of reward is calculated;
- (7) the agent's disutility of effort is calculated;
- (8) the agent's satisfaction is a function of effort, performance, act of nature, intrinsic reward, perceived equity of reward, compensation, and disutility of effort;
- (9) the principal's satisfaction is a function of output and compensation; and
- (10) the total satisfaction is the sum of the satisfactions of the agent and the principal.

The functions used in inference, selection of effort by the agent, and calculation of satisfaction are given below. The variables are multiplied by coefficients which denote an arbitrary priority of these variables for decision-making. Any such priority scheme may be used, or the functions replaced by knowledge-bases which help decide in selecting or calculating values for the decision variables. These functions are kept fixed for all the agents in the experiments. In function $f_1()$ for example, basic pay received the greatest weight, and terminal pay the least. Consideration of basic pay and share of output as the most important variables in determination of effort is consistent with the assumptions in the traditional principal-agent theory. Further, based on her experience of most agents, the principal expects the company environment and corporate ranking to play a more important role in the agent's acceptance of contracts and in

selection of effort than specific behavioral traits. This is reflected in the function $f_2()$, where the variable AT (Abilities and Traits) plays a vital role in the Porter and Lawler model in determining effort selection. The probability distribution of AT is derived from the probability distributions of the behavioral variables (excluding RISK, which plays a direct role in the model) as follows:

$$Pr[AT = i] = \frac{1}{9} \sum_{j=1, j \neq 4}^{10} Pr[b^{(j)} = i], \quad i = 1, \dots, 5,$$

where $b^{(j)}$ is the j -th behavioral variable, and $b^{(4)}$ is RISK (which is excluded).

The compensation variables enter the model in various ways, either directly as in effort selection or indirectly as in determination of intrinsic reward. Their major role is to induce the agent to select effort levels that lead to desired satisfaction levels.

The information available to the principal determines the weight of the different variables in the model and their contributory effect, or the derivation of AT from the behavioral variables. While the functions below reflect one such information system of the principal, others are possible.

$$f_1() = 13*BP + 12*S + 11*BO + 10*B + 9*SP + 8*TP;$$

$$f_2() = 7*CE + 6*WE + 5*ST + 4*AT;$$

$$\text{Effort} = g() = (f_1() + f_2() + 3*PPER + 2*IR + PEPR)/13;$$

$$\text{Output} = f_3() = \text{Effort} + \text{RISK};$$

$$\text{Performance, PERF} = f_4() = \text{Output} / \text{Effort};$$

$$\text{Intrinsic Reward, IR} = f_5() = (3*CE + 2*WE + ST)/6;$$

$$h_1() = 5*S + 4*BP + 3*BO + 2*SP + B;$$

$$PERR \equiv f_6() = (6*PERF + h_1())/21;$$

$$\text{Disutility of Effort, Dis} \equiv f_7() = -\text{Effort} / 10;$$

$$h_2() = 10*BP + 9*S + 8*BO + 7*SP + 6*B + 5*TP;$$

Satisfaction of the agent, S_{At} :

$$S_{At} \equiv f_8() = (12*PERF + 11*IR + h_2() + 4*PERR + 3*\text{Effort} - 2*RISK + \text{Dis})/66;$$

$$h_3() = BP + (S/10 * \text{Output}) + BO + TP + B + SP;$$

$$\text{Principal's Satisfaction, } S_{Pt} \equiv f_9 = \text{Output} - h_3();$$

$$\text{Total Satisfaction, } S_t \equiv f_{10}() = S_{At} + S_{Pt}.$$

9.3.4 Genetic Learning Details

Genetic learning by the principal requires a "fitness measure" for each rule. Here, the fitness of a rule is the (weighted) sum of the satisfactions of the principal and the agent, and normalized with respect to the full knowledge base. As already noted, the satisfaction of the principal is the utility of the principal's residuum, while the satisfaction of the agent is derived from the Porter-Lawler model of motivation. The average fitness of the knowledge base is derived, and the fitnesses of the individual rules are normalized to the interval [0,1]. One-point crossover and mutation are then applied to the knowledge base to yield the next generation of rules. A copy of the rule with the maximum fitness is passed unchanged to the next knowledge base. Pilot studies for this model showed that in no case did the maximum fitness across iterations peak after 200 iterations. Hence, 200 iterations were employed for all the experiments.

The three parameters which can be controlled in learning with genetic algorithms are the mating probability (MATE), the mutation probability (MUTATE), and the number of iterations (ITER). From trial simulations, a mating probability of 0.6, a mutation probability of 0.01, and 200 iterations for each run were deemed satisfactory, and were hence kept constant in all the experiments.

9.3.5 Statistics Captured for Analysis

The following statistics were collected for each simulation:

- (1) Average fitness of the principal's knowledge base;
- (2) Variance of knowledge base fitness;
- (3) Maximum fitness over all iterations of a run;
- (4) Entropy of fitnesses; and
- (5) Iteration when maximum fitness was first achieved.

These statistics are averaged across 10 runs for each experiment. The satisfaction index of the rules are normalized to the interval [0,1] to give fitness levels. Entropy is defined as the Shannon entropy, given by the formula

$$En\{f_i\} = - \sum_{i=1}^n f_i * \ln f_i,$$

where f_i is the fitness of the i -th rule in the knowledge base, and \ln is the natural logarithm. The maximum entropy possible is $\ln(\text{Number of Rules})$ and corresponds to the entropy of a distribution which occurs as the solution to a problem without any constraints (or information). In all the experiments, the possible maximum entropy is

therefore $\ln(501) \approx 6.2166061$. Addition of constraints or information (such as the value of the mean or variance) may result in a smaller entropy. The object of calculating the entropy of the knowledge base is to measure its informativeness. When the fitnesses, expressed as a distribution, achieve the maximum entropy while satisfying all the constraints of the system, the knowledge base is most informative yet maximally non-committal (see, for example, Jaynes 1982, 1986a, 1986b, 1991). An entropy value which is smaller than the maximum indicates some loss of information, while a larger entropy indicates unwarranted assumption of information. The entropy values will be compared across experiments to give an indication of the nature of the learned rules.

9.4 Results

The distribution of first iteration to achieve the maximum fitness bound is shown in Table 9.2 (expressed as a percentage) for the experiments. The table shows that there is more than a 38% chance of the maximum occurring within the first 30 iterations, a 50% chance of the maximum occurring within the first 60 iterations, and more than a 78% chance that it will do so within the first 120 iterations.

Learning appeared to converge quickly to the best knowledge base formed over the 200 learning episodes. Table 9.2 only indicates the way the learning process converges. Based on a number of pre-tests, this trend was found to be consistent. However, it should not be taken as an exact guide in any replication of the experiments.

Since random mutations in the learning process might result in rules which are not representative of the agent, the final knowledge base is processed to remove such

rules. The processing involves removal of those rules which have at least one antecedent value (i.e. value of a behavioral variable) which is not within one standard deviation range of the mean (given in Table 9.1). The processed knowledge bases of all the runs of each experiment are pooled to form the final knowledge base.

Table 9.3 shows the fitness statistics for the various experiments, where MATE = 0.6, MUTATION = 0.01 and ITER = 200 is fixed. Table 9.3 also shows the redundancy ratio of the knowledge base of each experiment. This is the ratio of the total number of rules to the number of distinct rules. This ratio may be greater than one because the learning process may generate copies of highly stable rules.

Table 9.4 shows the attained Shannon entropy of normalized fitness of the final knowledge base for each experiment. Table 9.4 also shows the theoretical maximum entropy of fitness (defined as the natural logarithm of the number of rules), and the ratio of the attained entropy to the maximum entropy. The fitness of each rule is multiplied by 10000 for readability.

Tables 9.5 - 9.34 summarize the results of the five experiments in detail. Tables 9.5, 9.11, 9.17, 9.23, and 9.29 show the frequency of values of the compensation variables in the final knowledge base. Tables 9.6, 9.12, 9.18, 9.24, and 9.30 show the range (minimum and maximum), mean, and standard deviation of the compensation variables. Tables 9.7, 9.13, 9.19, 9.25, and 9.31 show the results of Spearman correlation analysis on the final knowledge base. Tables 9.8, 9.9, 9.10, 9.14, 9.15, 9.16, 9.20, 9.21, 9.22, 9.26, 9.27, 9.28, 9.32, 9.33, and 9.34 deal with factor analysis of the final knowledge base. Tables 9.8, 9.14, 9.20, 9.26, and 9.32 list the eigenvalues

of the correlation matrix. Tables 9.9, 9.15, 9.21, 9.27, and 9.33 show the factor pattern of the direct solution (i.e. without rotation). Tables 9.10, 9.16, 9.22, 9.28, and 9.34 show the factor pattern of the varimax rotation.

The following rules having highest fitness in each experiment are displayed below for illustration (the rule representation format of Section 5.1 will be used; fitnesses, denoted FIT, are multiplied by 10,000 for convenience):

EXP 1: IF <3,2,5,2,3,4,4,4,3> THEN <4,1,1,1,1,1>;

EXP 2: IF <2,1,1,2,1,1,2,2,1> THEN <3,1,2,1,1,1>;

EXP 3: IF <1,4,3,4,3,3,5,5,4> THEN <5,1,3,1,1,1>;

EXP 4: IF <2,3,3,3,4,3,4,3,2> THEN <5,1,1,1,1,3>;

EXP 5: IF <3,4,3,4,4,5,5,5,4> THEN <5,1,2,1,1,3>.

9.5 Analysis of Results

In each of the experiments, letting the process run to completion usually improved the average fitness of the population, decreased its variance and increased its entropy. Several exceptions to this suggest that it may be a better strategy to store those knowledge bases generated during the learning process which possess desirable characteristics. Low variance indicates higher certainty, while higher entropy indicates a stable state close to a global optimum and uniformity in fitness for the rules of the population.

Agent #1 provides the maximum total satisfaction, followed in decreasing order by Agents #2, #5, #4, and #3 (Table 9.3). Interestingly, certain information did not

yield higher average fitness (or satisfaction) as can be seen by comparing Agents #5 and #1; a complete non-informative prior (as in the case of Agent #4) did not lead to lowest average fitness (the lowest was obtained by Agent #3). Furthermore, the result in this case is counter-intuitive when the behavioral characterizations of the different agents are considered. Agent #5 seems to be the best bet for this principal to maximize satisfaction. This is not the case. This result takes on added significance in view of the fact that Agent #5 faced an environment having low exogenous risk compared to that faced by Agent #1 (higher values for the risk variable in Table 9.1 denote less risk).

However, the uncertainty of the agent's performance in maximizing total satisfaction is least in the case of Agent #5 (about whom the principal has completely certain information), while it is the highest in the case of Agent #4 (about whom the principal has no information whatsoever). Agent #5 is followed in increasing order of uncertainty by Agents #3, #1, #2, and #4.

From Table 9.4, the ratio of the entropy of the normalized fitnesses of the knowledge base to the theoretical maximum gives an indication of how close the information content of the final knowledge base is to the theoretical maximum. It shows that the final knowledge base of the non-informative case (Agent #4) is least informative (while satisfying maximal non-committalness), while the case of certain information (Agent #5) shows a highly informative knowledge base. This is intuitively reasonable.

Tables 9.5, 9.6, 9.11, 9.12, 9.17, 9.18, 9.23, 9.24, 9.29 and 9.30 show the compensation recommendations for each of the five agents. The mean compensation value for each variable including the standard deviation from the mean helps in the task

of deciding on a specific compensation plan. For example, Agent #1 must be given high basic pay but as less of the other elements of compensation as possible, while Agent #2 should be given an above average (but not high) basic pay and a low amount of bonus (Table 9.11). Only in the non-informative case (Agent #4) a definite recommendation is made for the share of output to be as low as possible in the compensation plan offered to him (Table 9.24). Furthermore, if standard deviation from the mean compensation values is to be understood as the uncertainty regarding compensation, it is interesting to observe that in the case of Agent #4, the recommendations for compensation plans is more definitive than in the case of Agent #5 (as can be seen from comparing the standard deviations in Tables 9.24 and 9.30).

A few correlations at the 0.1 significance level among the compensation variables were observed (Tables 9.7, 9.13, 9.19, 9.25, and 9.31). For Agent #1, a mild positive correlation of 0.1173 was observed between Basic Pay and Terminal Pay (Table 9.7). For Agent #2, mild negative correlations between Basic Pay and Bonus (-0.2396) and between Basic Pay and Benefits (-0.1101) were observed. Bonus and Benefits were mildly positively correlated (Table 9.13). In the case of Agent #3, the following correlations were evident: Basic Pay and Share (-0.2124), Benefits and Share (0.2552), Bonus and Benefits (0.3042), and Benefits and Stock Participation (0.2762) (Table 9.19). No correlations were observed at all (at the 0.1 significance level) for Agent #4 (the non-informative case) (Table 9.25), while Agent #5 had the most number of significant correlations (7 out of a possible 15). However, all of these correlations were, without exception, very weak. Basic Pay formed weak negative correlations with Share (-0.0598)

and Terminal Pay (-0.0568), and weak positive correlations with Bonus (0.1591) and Stock Participation (0.0389). Benefits and Share were weakly positively correlated (0.0557). Stock Participation formed weak positive correlations with Basic Pay (0.0389), Bonus (0.0508) and Benefits (0.0605) (Table 9.31).

Without further research, the causes of these correlations cannot be known definitely. While the compensation schemes are definitely tailored to the behavioral characteristics of the agents, motivation theory does not enable (at the present state of the art) to make definitive causal connections between specific behavioral patterns and effort-inducing compensation. Directions for future research are described in Chapter 12.

Factor analysis of the final knowledge base of each experiment was carried out to see if the knowledge base had any significant factors. A factor with eigenvalue greater than one may be deemed to be significant since it accounts for more variation in the rules than any one variable alone. Table 9.9 provides a summary of pertinent data from Tables 9.8, 9.14, 9.20, 9.26, and 9.32.

The percentage of total variation accounted for by the significant factors is rather low, the maximum being for experiment 3. Experiment 4 required the maximum number of factors (almost as many as the number of variables, which is 16). Experiment 4 also had the highest average eigenvalue, and Experiment 5 the lowest. The number of significant factors was least in the case of Experiment 5, and each factor accounted for a greater proportion of the variation than the non-informative situation of Experiment 4.

This suggests that the final knowledge base of Experiment 4 is comparatively highly "fragmented" than that of Experiment 5.

Tables 9.9, 9.15, 9.21, 9.27, and 9.33 show the direct factor pattern. Variables that load high on a factor indicate a greater role played in explaining that factor. Moreover, each factor accounts for a small proportion of the total variation. A measure of the explanatory power of a variable may be the expected factor identification, defined as the sum of the products of each factor loading and the proportion of variation of that factor, the sum being taken over the total number of factors that account for all the variation in the population. Table 9.36 shows the expected factor identification of each of the compensation variables for each experiment. Table 9.37 shows the expected factor identification computed from the varimax rotated factor matrices.

Table 9.36 shows that except in Experiment 3, Basic Pay and Share did not have the highest explanatory measure. Comparing across all the five experiments and ranking the compensation variables, the following is the order of variables in decreasing explanatory measure:

- Benefits and stock participation (tied),
- Terminal pay,
- Bonus,
- Basic pay (also called Rent or Fixed pay), and
- Share.

A similar comparison from the data in Table 9.37 for the varimax rotated factors yields the following ordering of the compensation variables:

Benefits and stock participation (tied),
Basic pay and terminal pay (tied),
Share, and
Bonus.

Using the direct factor matrices, the expected factor identifications of behavioral variables were computed (Table 9.38). These variables were ranked and ordered across the 5 experiments. The exogenous risk variable, though not a behavioral variable, was included to study its relative importance also. The following is the decreasing order of explanatory power:

Experience,
Managerial skills,
General social skills,
Risk,
Physical qualities,
Communication skills,
Education,
Motivation,
Other personal skills, and
Age.

Using the varimax factor matrices, the expected factor identifications of behavioral variables were computed (Table 9.38). These variables were ranked and

ordered across the 5 experiments. The following is the decreasing order of explanatory power:

- Experience, risk, and physical qualities (tied),
- Managerial skills,
- Motivation,
- Age, and general social skills (tied),
- Education,
- Communication skills, and other personal skills (tied).

The above results and analysis support the hypothesis that behavioral characteristics and complex compensation plans play a significant role in determining good compensation rules (Hypothesis 1 in Sec. 9.2). However Hypothesis 2, regarding the high relative importance of Basic Pay and Share in the case of completely certain information (Experiment 5), has not been supported (see Sec. 9.2). The results show that even when complete and certain information is present, it is not reasonable for the principal to try to induce the agent to exert optimum effort by presenting a contract based solely on Basic Pay and Share of output. Further, the results provide a counterexample to the seemingly intuitive notion that either perfect information about the behavioral characteristics of the agent will yield the most satisfaction or that a complete lack of information about the agent will lead to minimum satisfaction. This suggests that Hypothesis 3 (see Sec. 9.2) is also not supported.

TABLE 9.1: Characterization of Agents

PERSONAL VARIABLE	EXP1 MEAN (SD)	EXP2 MEAN (SD)	EXP3 MEAN (SD)	EXP5 MEAN (SD)
COMPANY ENVIRONMENT, CE	3.40 (0.66)	3.60 (0.80)	4.50 (0.50)	5.00 (0.00)
WORK ENVIRONMENT, WE	3.60 (0.80)	3.60 (0.80)	4.60 (0.49)	5.00 (0.00)
STATUS INDEX, SI	4.20 (0.98)	4.20 (0.98)	4.70 (0.46)	5.00 (0.00)
ABILITIES AND TRAITS, AT	3.51 (1.16)	1.46 (0.74)	3.73 (1.18)	4.11 (0.74)
PROB. (EFFORT -> REWARD), PPER	2.90 (0.70)	2.90 (0.70)	4.60 (0.66)	4.00 (0.00)
BEHAVIORAL VARIABLE				
EXPERIENCE, X	3.65 (0.78)	1.40 (0.66)	1.40 (0.66)	3.00 (0.00)
EDUCATION, D	2.00 (0.00)	1.00 (0.00)	4.20 (0.40)	4.00 (0.00)
AGE, A	5.00 (0.00)	1.00 (0.00)	3.00 (0.00)	3.00 (0.00)
RISK	2.10 (1.38)	2.10 (1.38)	3.90 (0.94)	4.00 (0.00)
GENERAL SOCIAL SKILLS, GSS	3.00 (1.55)	1.70 (0.78)	3.80 (0.87)	4.00 (0.00)
MANAGERIAL SKILLS, OMS	4.05 (0.67)	1.40 (0.66)	3.40 (0.92)	5.00 (0.00)
MOTIVATION, M	3.60 (0.66)	1.50 (0.50)	4.90 (0.30)	5.00 (0.00)
PHYSICAL QUALITIES, PQ	3.60 (1.28)	2.3 (1.1)	4.60 (0.49)	5.00 (0.00)
COMMUNICATION SKILLS, L	3.95 (0.59)	1.50 (0.67)	4.30 (0.64)	4.00 (0.00)
OTHERS, OPC	2.77 (0.61)	1.30 (0.64)	4.00 (0.78)	4.00 (0.00)

TABLE 9.2: Iteration of First Occurrence of Maximum Fitness

RANGE	PERCENTAGE	RANGE	PERCENTAGE
[1,10]	14	(100,110]	4
(10,20]	14	(110,120]	4
(20,30]	10	(120,130]	0
(30,40]	6	(130,140]	4
(40,50]	2	(140,150]	2
(50,60]	4	(150,160]	2
(60,70]	6	(160,170]	2
(70,80]	4	(170,180]	4
(80,90]	4	(180,190]	0
(90,100]	6	(190,200]	6

TABLE 9.3: Learning Statistics for Fitness of Final Knowledge Bases

Experiment	Number of Rules	Redundancy Ratio	Minimum	Maximum	Mean	S.D.
1	199	1.3724	13.96	27.19	20.29	2.87
2	397	1.3690	7.77	27.16	19.92	3.18
3	63	1.1455	11.09	24.66	19.71	2.42
4	74	1.1563	11.92	26.68	19.82	3.69
5	1965	5.7794	3.09	24.72	19.94	2.35

TABLE 9.4: Entropy of Final Knowledge Bases and Closeness to the Maximum

Experiment	Number of Rules	Entropy	Maximum Entropy	Ratio ¹
1	199	5.2834	5.2933	0.9981
2	397	5.9709	5.9839	0.9978
3	63	4.1355	4.1431	0.9982
4	74	4.2869	4.3041	0.9960
5	1965	7.5760	7.5833	0.9990

¹Ratio of Entropy to Maximum Entropy

TABLE 9.5: Frequency (as Percentage) of Values of Compensation Variables in the Final Knowledge Base in Experiment 1

Compensation Variables	Values of the Variable				
	1	2	3	4	5
Basic Pay	3.0	4.5	13.1	38.2	41.2
Share	97.5	2.0	0.5	0.0	0.0
Bonus	61.8	22.6	5.5	5.5	4.5
Terminal Pay	93.0	2.0	3.0	1.5	0.5
Benefits	82.9	9.5	1.5	3.0	3.0
Stock Participation	74.4	18.1	6.0	1.5	0.0

TABLE 9.6: Range, Mean and Standard Deviation of Values of Compensation Variables in the Final Knowledge Base in Experiment 1

Variable	Minimum	Maximum	Mean	S.D.
BP	1.00	5.00	4.1005025	0.9949112
S	1.00	3.00	1.0301508	0.1987219
BO	1.00	5.00	1.6834171	1.0987947
TP	1.00	5.00	1.1457286	0.5807252
B	1.00	5.00	1.3366834	0.8945464
SP	1.00	4.00	1.3467337	0.6631551

TABLE 9.7: Correlation Analysis of Values of Compensation Variables in the Final Knowledge Base in Experiment 1
(Spearman Correlation Coefficients in the first row for each variable,
Prob > |R| under Ho: Rho=0 in the second)

	BP	S	BO	TP	B	SP
BP	1.00000	0.04942	0.11030	0.11728	-0.07280	0.02989
	0.0	0.4882	0.1209	0.0990	0.3068	0.6752
S	0.04942	1.00000	0.03915	-0.04413	0.00558	-0.09337
	0.4882	0.0	0.5830	0.5360	0.9377	0.1896
BO	0.11030	0.03915	1.00000	-0.03333	0.00579	-0.03130
	0.1209	0.5830	0.0	0.6402	0.9354	0.6607
TP	0.11728	-0.04413	-0.03333	1.00000	0.02864	-0.00110
	0.0990	0.5360	0.6402	0.0	0.6880	0.9877
B	-0.07280	0.00558	0.00579	0.02864	1.00000	-0.04710
	0.3068	0.9377	0.9354	0.6880	0.0	0.5089
SP	0.02989	-0.09337	-0.03130	-0.00110	-0.04710	1.00000
	0.6752	0.1896	0.6607	0.9877	0.5089	0.0

TABLE 9.8: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 1
(Eigenvalues of the Correlation Matrix)
Total = 11 Average = 0.6875

	1	2	3	4	5	6
Eigenvalue	1.494645	1.352410	1.160236	1.116413	1.051618	1.012871
Difference	0.142235	0.192174	0.043823	0.064795	0.038747	0.080984
Proportion	0.1359	0.1229	0.1055	0.1015	0.0956	0.0921
Cumulative	0.1359	0.2588	0.3643	0.4658	0.5614	0.6535
	7	8	9	10	11	12
Eigenvalue	0.931887	0.823970	0.757124	0.714118	0.584709	0.000000
Difference	0.107916	0.066847	0.043006	0.129409	0.584709	0.000000
Proportion	0.0847	0.0749	0.0688	0.0649	0.0532	0.0000
Cumulative	0.7382	0.8131	0.8819	0.9468	1.0000	1.0000
	13	14	15	16		
Eigenvalue	0.000000	0.000000	0.000000	0.000000		
Difference	0.000000	0.000000	0.000000	0.000000		
Proportion	0.0000	0.0000	0.0000	0.0000		
Cumulative	1.0000	1.0000	1.0000	1.0000		

TABLE 9.9: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 1 - Factor Pattern

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6
X	-0.38741	-0.32701	-0.33959	0.31473	-0.02677	-0.10644
D	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
A	-0.00000	0.00000	0.00000	0.00000	0.00000	-0.00000
RISK	-0.39646	-0.13369	0.50581	0.25216	-0.11529	0.46880
GSS	0.17684	0.23305	0.28615	-0.36547	-0.61216	0.30750
OMS	-0.00000	0.00000	0.00000	-0.00000	-0.00000	0.00000
M	0.45141	0.44846	-0.30819	0.01878	-0.02665	-0.04173
PQ	0.54728	-0.53206	0.16099	0.05279	0.21852	0.00511
L	-0.00000	-0.00000	0.00000	-0.00000	-0.00000	0.00000
OPC	0.00000	-0.00000	-0.00000	0.00000	0.00000	-0.00000
BP	0.24127	0.19271	0.26919	0.66974	0.26082	0.24682
S	0.15889	0.06246	-0.59932	0.06671	0.04408	0.52285
BO	0.55605	-0.44006	0.19261	0.03879	-0.21608	-0.29675
TP	0.28107	0.47123	0.27576	-0.15331	0.47397	0.01654
B	-0.31708	-0.16489	0.12134	-0.52885	0.51514	0.06433
SP	-0.28396	0.45292	0.16366	0.24367	-0.08786	-0.50860
	Factor 7	Factor 8	Factor 9	Factor 10	Factor 11	
X	0.53275	0.32110	0.18254	-0.24025	0.19642	
D	0.00000	0.00000	0.00000	0.00000	0.00000	
A	0.00000	-0.00000	-0.00000	0.00000	0.00000	
RISK	0.15023	0.09305	0.00794	0.48441	0.08071	
GSS	0.13197	0.30346	0.03435	-0.34316	-0.03492	
OMS	-0.00000	0.00000	0.00000	-0.00000	0.00000	
M	0.38181	0.20000	-0.44346	0.31541	0.12412	
PQ	0.21527	0.25601	0.02163	0.04990	-0.47547	
L	-0.00000	0.00000	0.00000	-0.00000	0.00000	
OPC	0.00000	-0.00000	-0.00000	0.00000	-0.00000	
BP	-0.17272	0.06504	-0.24630	-0.38664	0.10238	
S	-0.38033	0.29019	0.29471	0.12585	-0.01855	
BO	-0.26496	0.16193	0.12058	0.12813	0.44320	
TP	0.28937	-0.02725	0.51720	0.01514	0.14920	
B	-0.15392	0.41618	-0.28532	-0.07102	0.15815	
SP	-0.25267	0.47536	0.12030	0.12243	-0.20590	

Notes: Final Communality Estimates total 11.0 and are as follows: 0.0 for D, A, OMS, L, and OPC; 1.0 for the rest of the variables.

TABLE 9.10: Experiment 1 - Varimax Rotation

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6
X	0.00592	-0.00592	-0.06160	-0.00152	0.03289	-0.03681
D	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
A	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
RISK	0.00697	0.01680	0.02680	-0.04246	0.99219	0.03829
GSS	-0.00519	-0.02315	0.99595	-0.00698	0.02638	-0.02607
OMS	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
M	0.00646	-0.06592	0.03811	0.04031	-0.08515	0.02529
PQ	-0.10498	0.00895	-0.01482	-0.00856	-0.01218	0.03838
L	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
OPC	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
BP	0.02878	-0.05624	-0.02624	0.01905	0.03809	0.99414
S	-0.04050	-0.01274	-0.00694	0.99688	-0.04193	0.01893
BO	-0.02424	-0.05007	0.01906	-0.01453	-0.04793	0.00752
TP	0.01929	0.00864	0.02160	-0.02133	-0.02683	0.04298
B	-0.00484	0.99454	-0.02324	-0.01282	0.01670	-0.05616
SP	0.99305	-0.00485	-0.00525	-0.04099	0.00693	0.02892
	Factor 7	Factor 8	Factor 9	Factor 10	Factor 11	
X	-0.07567	0.99216	-0.05006	-0.03047	0.01219	
D	0.00000	0.00000	0.00000	0.00000	0.00000	
A	0.00000	0.00000	0.00000	0.00000	0.00000	
RISK	-0.02644	0.03284	-0.04665	-0.08451	-0.01179	
GSS	0.02122	-0.06089	0.01826	0.03736	-0.01414	
OMS	0.00000	0.00000	0.00000	0.00000	0.00000	
M	0.06591	-0.03072	-0.01159	0.98940	0.01758	
PQ	0.02316	0.01277	0.15553	0.01808	0.98070	
L	0.00000	0.00000	0.00000	0.00000	0.00000	
OPC	0.00000	0.00000	0.00000	0.00000	0.00000	
BP	0.04274	-0.03658	0.00741	0.02495	0.03682	
S	-0.02106	-0.00149	-0.01395	0.03942	-0.00818	
BO	-0.03591	-0.05154	0.98287	-0.01188	0.15448	
TP	0.99215	-0.07568	-0.03483	0.06528	0.02213	
B	0.00860	-0.00589	-0.04828	-0.06490	0.00843	
SP	0.01928	0.00591	-0.02373	0.00641	-0.10112	

Notes: Final Communality Estimates total 11.0 and are as follows: 0.0 for D, A, OMS, L, and OPC; 1.0 for the rest of the variables.

TABLE 9.11: Frequency (as Percentage) of Values of Compensation Variables in the Final Knowledge Base in Experiment 2

Compensation Variable	VALUES OF THE VARIABLE				
	1	2	3	4	5
Basic Pay	6.5	2.0	17.1	45.3	29.0
Share	95.7	1.8	0.8	1.3	0.5
Bonus	50.1	22.7	7.6	13.4	6.3
Terminal Pay	93.7	3.3	1.3	0.5	1.3
Benefits	85.1	8.6	3.5	2.0	0.8
Stock	87.9	6.8	2.0	1.8	1.5

TABLE 9.12: Range, Mean and Standard Deviation of Values of Compensation Variables in the Final Knowledge Base in Experiment 2

Variable	Minimum	Maximum	Mean	S.D.
BP	1.00	5.00	3.8816121	1.0582000
S	1.00	5.00	1.0906801	0.4839221
BO	1.00	5.00	2.0302267	1.2964961
TP	1.00	5.00	1.1234257	0.5617257
B	1.00	5.00	1.2468514	0.6849916
SP	1.00	5.00	1.2216625	0.7079878

TABLE 9.13: Correlation Analysis of Values of Compensation Variables in the Final Knowledge Base in Experiment 2

(Spearman Correlation Coefficients in the first row for each variable, Prob> |R| under Ho: $\rho=0$ in the second)

	BP	S	BO	TP	B	SP
BP	1.00000	0.02951	-0.23955	0.05064	-0.11008	0.01298
	0.0	0.5578	0.0001	0.3142	0.0283	0.7965
S	0.02951	1.00000	0.03275	-0.00414	0.06030	0.00038
	0.5578	0.0	0.5153	0.9344	0.2307	0.9940
BO	-0.23955	0.03275	1.00000	0.01020	0.10281	-0.02808
	0.0001	0.5153	0.0	0.8394	0.0406	0.5770
TP	0.05064	-0.00414	0.01020	1.00000	0.04402	-0.00848
	0.3142	0.9344	0.8394	0.0	0.3817	0.8663
B	-0.11008	0.06030	0.10281	0.04402	1.00000	0.01402
	0.0283	0.2307	0.0406	0.3817	0.0	0.7807
SP	0.01298	0.00038	-0.02808	-0.00848	0.01402	1.00000
	0.7965	0.9940	0.5770	0.8663	0.7807	0.0

TABLE 9.14: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 2 - Eigenvalues of the Correlation Matrix

Factors	1	2	3	4	5	6
Eigenvalue	1.562150	1.349480	1.288563	1.186437	1.075113	1.008861
Difference	0.212669	0.060917	0.102126	0.111324	0.066252	0.039300
Proportion	0.1202	0.1038	0.0991	0.0913	0.0827	0.0776
Cumulative	0.1202	0.2240	0.3231	0.4144	0.4971	0.5747
	7	8	9	10	11	12
Eigenvalue	0.969560	0.913091	0.869975	0.797047	0.744512	0.637679
Difference	0.056469	0.043117	0.072927	0.052535	0.106833	0.040147
Proportion	0.0746	0.0702	0.0669	0.0613	0.0573	0.0491
Cumulative	0.6492	0.7195	0.7864	0.8477	0.9050	0.9540
	13	14	15	16		
Eigenvalue	0.597532	0.000000	0.000000	0.000000		
Difference	0.597532	0.000000	0.000000			
Proportion	0.0460	0.0000	0.0000	0.0000		
Cumulative	1.0000	1.0000	1.0000	1.0000		

TABLE 9.15: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 2 - Factor Pattern

Factor	1	2	3	4	5	6	7
X	0.23470	0.40822	0.41989	0.10143	-0.06455	-0.47522	-0.13029
D	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
A	-0.00000	-0.00000	0.00000	-0.00000	0.00000	0.00000	-0.00000
RISK	0.40696	0.10345	-0.15114	-0.43374	0.06063	0.00857	0.09780
GSS	0.71945	0.12878	0.06651	0.25299	-0.09613	-0.02708	0.21848
OMS	0.15988	0.26111	-0.23533	0.55088	0.53138	0.02406	-0.24275
M	0.52994	-0.06812	0.35756	0.04245	0.12138	-0.00603	0.43643
PQ	-0.49072	0.18271	0.19086	0.25694	0.29339	-0.29651	0.18784
L	-0.43182	0.11417	0.46917	-0.00994	-0.26917	-0.20723	0.25407
OPC	-0.00000	-0.00000	-0.00000	0.00000	0.00000	-0.00000	-0.00000
BP	0.00206	0.73317	0.02383	-0.14116	0.09366	0.15727	-0.00144
S	0.15005	0.08440	0.48987	0.00873	-0.35507	0.37735	-0.43373
BO	0.16081	-0.64356	0.15230	0.08186	0.15639	-0.23750	0.04318
TP	-0.11221	0.09559	0.22871	-0.50489	0.50146	0.27257	0.26726
B	-0.07398	-0.24228	0.54532	0.23797	0.32011	0.41360	-0.14065
SP	-0.15402	0.09822	-0.17831	0.46301	-0.29859	0.42643	0.51464
Factor	8	9	10	11	12	13	
X	-0.34093	0.12914	0.28987	-0.21169	-0.06036	-0.28162	
D	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	
A	0.00000	-0.00000	0.00000	0.00000	-0.00000	0.00000	
RISK	0.64290	0.27270	0.17846	-0.21469	-0.03416	-0.18056	
GSS	0.03558	-0.12787	0.18071	0.00686	-0.21839	0.49158	
OMS	0.14111	0.00621	0.04899	-0.09147	0.41684	0.03269	
M	-0.03356	-0.10222	-0.53502	0.03731	0.16803	-0.22844	
PQ	0.21954	0.40686	-0.28804	-0.04598	-0.29716	0.16419	
L	0.35726	-0.28314	0.19616	0.00520	0.37074	0.12874	
OPC	-0.00000	0.00000	0.00000	0.00000	0.00000	-0.00000	
BP	0.05725	-0.04855	0.03755	0.62106	-0.07920	-0.09706	
S	0.05131	0.43120	-0.18876	-0.00933	0.15673	0.15771	
BO	0.00810	0.35568	0.28456	0.47545	0.11527	-0.02132	
TP	-0.35860	0.14954	0.18191	-0.15212	0.14461	0.21390	
B	0.20155	-0.28888	0.19343	-0.05661	-0.30269	-0.17940	
SP	-0.10104	0.29287	0.22341	-0.05877	0.02784	-0.18570	

Notes: Final Communality Estimates total 13.0 and are as follows: 0.0 for D, A, and OPC; 1.0 for the rest of the variables.

TABLE 9.16: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 2 -Varimax Rotated Factor Pattern

Factor	1	2	3	4	5	6	7
X	0.03492	0.99126	0.02196	-0.01687	0.01069	0.01925	-0.03964
E	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
A	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
RISK	0.02987	-0.01676	-0.04298	0.99302	0.02126	-0.00162	-0.04582
GSS	0.12998	0.08431	-0.09124	0.06914	-0.05572	0.05292	0.02567
OMS	-0.00638	0.01924	0.04137	-0.00161	-0.03958	0.99160	0.01619
M	0.98841	0.03512	-0.02218	0.03018	0.02264	-0.00653	-0.02308
PQ	-0.02223	0.02208	0.98938	-0.04343	0.01497	0.04168	0.02901
L	-0.01908	0.03370	0.08208	-0.03191	-0.00782	-0.07776	0.01142
OPC	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
BP	-0.01095	0.04783	0.02059	0.02922	0.04601	0.03992	0.00344
S	0.02705	0.05363	-0.01922	0.00868	-0.01382	-0.03522	0.00210
BO	0.03723	-0.01392	-0.00442	0.00168	-0.00869	-0.02235	-0.02816
TP	0.02208	0.01055	0.01474	0.02112	0.99503	-0.03917	-0.02333
B	0.03185	-0.01912	0.01580	-0.04609	0.03829	0.03085	-0.01322
SP	-0.02242	-0.03903	0.02838	-0.04535	-0.02323	0.01596	0.99633
Factor	8	9	10	11	12	13	
X	0.05425	-0.01391	-0.01919	0.03347	0.04754	0.08050	
D	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	
A	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	
RISK	0.00873	0.00165	-0.04594	-0.03152	0.02887	0.06572	
GSS	0.01523	0.01885	0.00031	-0.05231	0.01946	0.97603	
OMS	-0.03559	-0.02242	0.03091	-0.07707	0.03965	0.05044	
M	0.02757	0.03741	0.03224	-0.01900	-0.01107	0.12505	
PQ	-0.01954	-0.00444	0.01591	0.08190	0.02059	-0.08759	
L	0.00709	-0.02679	0.05532	0.98880	0.02057	-0.05035	
OPC	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	
BP	0.02113	-0.10936	-0.03162	0.02050	0.98915	0.01871	
S	0.99514	-0.00501	0.05916	0.00695	0.02076	0.01446	
BO	-0.00505	0.99110	0.04331	-0.02656	-0.10871	0.01810	
TP	-0.01382	-0.00867	0.03796	-0.00765	0.04522	-0.05255	
B	0.05972	0.04336	0.99208	0.05475	-0.03138	0.00032	
SP	0.00209	-0.02757	-0.01314	0.01116	0.00340	0.02406	

Notes: Final Communality Estimates total 11.0 and are as follows: 0.0 for D, A, and OPC; 1.0 for the rest of the variables.

TABLE 9.17: Frequency (as Percentage) of Values of Compensation Variables in the Final Knowledge Base in Experiment 3

COMPENSATION VARIABLE	VALUES OF THE VARIABLE				
	1	2	3	4	5
BASIC PAY	11.1	15.9	14.3	17.5	41.3
SHARE	92.1	6.3	0.0	1.6	0.0
BONUS	30.2	33.3	25.4	7.9	3.2
TERMINAL PAY	85.7	6.3	3.2	1.6	3.2
BENEFITS	76.2	12.7	3.2	3.2	4.8
STOCK	66.7	14.3	6.3	9.5	3.2

TABLE 9.18: Range, Mean and Standard Deviation of Values of Compensation Variables in the Final Knowledge Base in Experiment 3

Variable	Minimum	Maximum	Mean	S.D.
BP	1.00	5.00	3.6190476	1.4416452
S	1.00	4.00	1.1111111	0.4439962
BO	1.00	5.00	2.2063492	1.0649660
TP	1.00	5.00	1.3015873	0.8731648
B	1.00	5.00	1.4761905	1.0450674
SP	1.00	5.00	1.6825397	1.1475837

TABLE 9.19: Correlation Analysis of Values of Compensation Variables in the Final Knowledge Base in Experiment 3

	BP	S	BO	TP	B	SP
BP	1.00000	-0.21239	0.04193	0.09919	-0.06653	0.16112
	0.0	0.0947	0.7442	0.4392	0.6044	0.2071
S	-0.21239	1.00000	0.13992	0.06965	0.25522	-0.06207
	0.0947	0.0	0.2741	0.5875	0.0435	0.6289
BO	0.04193	0.13992	1.00000	-0.02696	0.30417	0.02454
	0.7442	0.2741	0.0	0.8339	0.0154	0.8486
TP	0.09919	0.06965	-0.02696	1.00000	-0.05317	-0.09539
	0.4392	0.5875	0.8339	0.0	0.6790	0.4571
B	-0.06653	0.25522	0.30417	-0.05317	1.00000	0.27619
	0.6044	0.0435	0.0154	0.6790	0.0	0.0284
SP	0.16112	-0.06207	0.02454	-0.09539	0.27619	1.00000
	0.2071	0.6289	0.8486	0.4571	0.0284	0.0

TABLE 9.20: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 3 - Eigenvalues of the Correlation Matrix

Factor	1	2	3	4	5	6
Eigenvalue	2.051970	1.485699	1.393947	1.302750	1.019307	0.766105
Difference	0.566272	0.091752	0.091196	0.283444	0.253202	0.134694
Proportion	0.2052	0.1486	0.1394	0.1303	0.1019	0.0766
Cumulative	0.2052	0.3538	0.4932	0.6234	0.7254	0.8020
Factor	7	8	9	10	11	12
Eigenvalue	0.631411	0.529438	0.495105	0.324267	0.0000	0.0000
Difference	0.101973	0.034334	0.170837	0.324267	0.0000	0.0000
Proportion	0.0631	0.0529	0.0495	0.0324	0.0000	0.0000
Cumulative	0.8651	0.9181	0.9676	1.0000	1.0000	1.0000
Factor	13	14	15	16		
Eigenvalue	0.0000	0.0000	0.0000	0.0000		
Difference	0.0000	0.0000	0.0000			
Proportion	0.0000	0.0000	0.0000	0.0000		
Cumulative	1.0000	1.0000	1.0000	1.0000		

TABLE 9.21: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 3 - Factor Pattern

FACTOR	1	2	3	4	5
A	-0.59074	-0.06170	-0.06894	0.46822	-0.20218
D	0.00000	0.00000	0.00000	0.00000	0.00000
A	-0.00000	0.00000	0.00000	-0.00000	-0.00000
RISK	0.10996	0.76295	0.33072	-0.11407	-0.01594
GSS	0.85184	0.10329	0.21497	-0.12037	-0.04482
OMS	0.80467	0.01491	0.07961	0.16731	-0.18963
M	-0.00000	0.00000	-0.00000	0.00000	0.00000
PQ	-0.00000	0.00000	-0.00000	0.00000	0.00000
L	-0.00000	0.00000	-0.00000	0.00000	0.00000
OPC	-0.00000	0.00000	-0.00000	0.00000	0.00000
BP	-0.39157	0.65888	0.22137	0.06324	0.23306
S	0.17892	-0.38179	0.38789	0.52832	0.35570
BO	0.13728	0.13920	-0.35713	-0.07526	0.82400
TP	-0.12358	-0.02483	0.78267	0.33425	0.10374
B	0.29624	0.09962	-0.43582	0.64893	0.08771
SP	0.10276	0.52831	-0.31264	0.45432	-0.24887
FACTOR	6	7	8	9	10
X	0.50944	0.05741	-0.05787	0.29569	0.16957
D	0.00000	0.00000	0.00000	0.00000	0.00000
A	-0.00000	0.00000	-0.00000	0.00000	-0.00000
RISK	0.19694	-0.09151	-0.48027	-0.04820	-0.05509
GSS	0.03364	-0.02060	0.09967	0.10779	0.42176
OMS	0.29181	0.03604	0.17912	0.22710	-0.33448
M	-0.00000	0.00000	-0.00000	-0.00000	0.00000
PQ	-0.00000	0.00000	0.00000	-0.00000	0.00000
L	-0.00000	0.00000	-0.00000	-0.00000	0.00000
OPC	-0.00000	0.00000	0.00000	0.00000	0.00000
BP	-0.19754	-0.26443	0.35986	0.25771	-0.01930
S	-0.35026	-0.00881	-0.28183	0.25181	-0.02289
BO	0.27700	0.26878	0.01950	0.01442	0.00576
TP	0.18578	0.19317	0.20941	-0.36516	0.00462
B	0.06503	-0.44763	0.00951	-0.27895	0.03276
SP	-0.32300	0.48794	0.01272	-0.03113	0.02641

Notes: Final Communality Estimates total 10.0 and are as follows: 0.0 for D, A, M, PQ, L, and OPC; 1.0 for the rest of the variables.

TABLE 9.22: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 3 - Varimax Rotated Factor Pattern

Factor	1	2	3	4	5
X	0.96928	-0.08142	-0.03918	0.03545	0.04428
D	0.00000	0.00000	0.00000	0.00000	0.00000
A	0.00000	0.00000	0.00000	0.00000	0.00000
RISK	-0.04055	0.04253	0.97181	0.16514	-0.02720
GSS	-0.25455	0.32873	0.10890	-0.08646	0.03277
OMS	-0.08603	0.93594	0.04323	-0.13317	0.11829
M	0.00000	0.00000	0.00000	0.00000	0.00000
PQ	0.00000	0.00000	0.00000	0.00000	0.00000
L	0.00000	0.00000	0.00000	0.00000	0.00000
OPC	0.00000	0.00000	0.00000	0.00000	0.00000
BP	0.03686	-0.12283	0.16709	0.96833	-0.01802
S	-0.01604	0.03958	-0.08037	-0.03568	0.06811
BO	-0.05005	0.00008	0.00895	0.01872	0.06265
TP	0.07511	0.01390	0.06546	0.07241	-0.06475
B	0.04250	0.10437	-0.02704	-0.01793	0.97700
SP	0.02553	0.04427	0.07087	0.06923	0.13070
Factor	6	7	8	9	10
X	-0.01562	0.02692	-0.05402	0.07670	-0.19835
D	0.00000	0.00000	0.00000	0.00000	0.00000
A	0.00000	0.00000	0.00000	0.00000	0.00000
RISK	-0.08227	0.07314	0.00945	0.06551	0.08735
GSS	0.04743	0.00360	0.01355	0.00660	0.89680
OMS	0.04351	0.05169	-0.00050	0.01576	0.27965
M	0.00000	0.00000	0.00000	0.00000	0.00000
PQ	0.00000	0.00000	0.00000	0.00000	0.00000
L	0.00000	0.00000	0.00000	0.00000	0.00000
OPC	0.00000	0.00000	0.00000	0.00000	0.00000
BP	-0.03644	0.07236	0.02019	0.07353	-0.07257
S	0.98027	-0.02612	-0.00494	0.15086	0.03751
BO	-0.00482	-0.00515	0.99439	-0.06454	0.01048
TP	0.15371	-0.04981	-0.06857	0.97448	0.00496
B	0.06843	0.13345	0.06613	-0.06398	0.02754
SP	-0.02605	0.98358	-0.00542	-0.04843	0.00392

Notes: Final Communality Estimates total 10.0 and are as follows: 0.0 for D, A, M, PQ, L, and OPC; 1.0 for the rest of the variables.

TABLE 9.23: Frequency (as Percentage) of Values of Compensation Variables in the Final Knowledge Base in Experiment 4

COMPENSATION VARIABLE	VALUES OF THE VARIABLE				
	1	2	3	4	5
BASIC PAY	8.0	10.8	6.8	20.3	54.1
SHARE	100.0	0.0	0.0	0.0	0.0
BONUS	62.2	24.3	4.1	8.1	1.4
TERMINAL PAY	82.4	5.4	5.4	4.1	2.7
BENEFITS	78.4	12.2	6.8	1.4	1.4
STOCK	82.4	14.9	1.4	0.0	1.4

TABLE 9.24: Range, Mean and Standard Deviation of Values of Compensation Variables in the Final Knowledge Base in Experiment 4

Variable	Minimum	Maximum	Mean	S.D.
BP	1.0000	5.0000	4.0135135	1.3395281
S	1.0000	1.0000	1.0000000	0
BO	1.0000	5.0000	1.6216216	0.9890178
TP	1.0000	5.0000	1.3918919	0.9625532
B	1.0000	5.0000	1.3513514	0.7839561
SP	1.0000	5.0000	1.2297297	0.6092281

TABLE 9.25: Correlation Analysis of Values of Compensation Variables in the Final Knowledge Base in Experiment 4

(Spearman Correlation Coefficients in the first row for each variable,
Prob > |R| under Ho: Rho=0 in the second)

	BP	S	BO	TP	B	SP
BP	1.00000	.	-0.07927	-0.12735	0.17890	0.09947
	0.00000		0.5020	0.2796	0.1272	0.3991
S	.	1.00000
		0.00000				
BO	-0.07927	.	1.00000	0.04158	-0.05059	-0.05058
	0.5020		0.00000	0.7250	0.6686	0.6687
TP	-0.12735	.	0.04158	1.00000	-0.15591	-0.03370
	0.2796		0.7250	0.00000	0.1847	0.7756
B	0.17890	.	-0.05059	-0.15591	1.00000	-0.07384
	0.1272		0.6686	0.1847	0.00000	0.5318
SP	0.09947	.	-0.05058	-0.03370	-0.07384	1.00000
	0.3991		0.6687	0.7756	0.5318	0.0

TABLE 9.26: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 4 - Eigenvalues of the Correlation Matrix
Total = 15 Average = 0.9375

Factor	1	2	3	4	5	6
Eigenvalue	2.266645	1.820044	1.740554	1.392479	1.222659	1.127880
Difference	0.446601	0.079490	0.348075	0.169820	0.094779	0.081301
Proportion	0.1511	0.1213	0.1160	0.0928	0.0815	0.0752
Cumulative	0.1511	0.2724	0.3885	0.4813	0.5628	0.6380
Factor	7	8	9	10	11	12
Eigenvalue	1.046579	0.911929	0.720692	0.673039	0.590800	0.540745
Difference	0.134650	0.191237	0.047653	0.082239	0.050055	0.139484
Proportion	0.0698	0.0608	0.0480	0.0449	0.0394	0.0360
Cumulative	0.7078	0.7686	0.8166	0.8615	0.9009	0.9369
Factor	13	14	15	16		
Eigenvalue	0.401261	0.330169	0.214527	0.000000		
Difference	0.071092	0.115642	0.214527			
Proportion	0.0268	0.0220	0.0143	0.0000		
Cumulative	0.9637	0.9857	1.0000	1.0000		

TABLE 9.27: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 4 - Factor Pattern

Factor	1	2	3	4	5
X	0.39987	0.52286	-0.38231	0.31635	0.00739
D	0.65507	0.02692	0.14216	-0.45863	0.14832
A	-0.69786	-0.01735	-0.21577	-0.09847	-0.45496
RISK	-0.56345	0.48709	0.00618	0.07723	0.29746
GSS	0.12282	0.34586	0.57027	0.25255	-0.09610
OMS	-0.18220	-0.28764	0.47685	0.37232	-0.13876
M	-0.10416	0.01902	0.71620	0.06443	0.17785
PQ	0.05916	0.70795	-0.11360	0.13412	0.32138
L	-0.26565	0.09598	0.42751	-0.65510	0.17288
OPC	0.16943	0.49388	-0.14315	-0.39151	-0.43841
BP	0.66360	-0.28499	0.05754	0.13749	-0.12937
S	-0.00000	0.00000	-0.00000	0.00000	-0.00000
BO	-0.09702	-0.33988	-0.32518	-0.18613	0.66939
TP	-0.19261	-0.15261	-0.17228	0.42969	0.19808
B	0.47077	-0.02466	0.13058	0.11620	0.10372
SP	-0.05247	0.36153	0.30017	0.08792	0.06952
Factor	6	7	8	9	10
X	0.11952	-0.10795	0.17609	-0.14149	-0.19329
D	-0.15335	-0.24955	0.24565	0.01985	0.09769
A	-0.06005	0.00624	-0.11661	0.29239	0.18458
RISK	0.39016	-0.07852	-0.10355	-0.06561	-0.11007
GSS	0.13077	0.31075	0.14084	-0.09173	0.39591
OMS	0.30435	-0.27467	0.39357	-0.17861	0.10779
M	-0.27035	0.15311	-0.03593	0.09622	-0.51882
PQ	-0.25059	0.18421	-0.20210	-0.05302	0.22436
L	-0.03394	0.21155	0.01226	-0.07221	0.14553
OPC	0.03602	0.11633	0.44300	0.11891	-0.14114
BP	-0.15691	0.01890	-0.35834	-0.07636	0.12265
S	0.00000	-0.00000	0.00000	0.00000	-0.00000
BO	0.19577	-0.06597	0.22588	0.05967	0.12389
TP	-0.53566	0.27996	0.43910	0.23157	0.07421
B	0.52889	0.24791	-0.08326	0.59940	-0.00482
SP	-0.23789	-0.73091	-0.06188	0.34125	0.12271
Factor	11	12	13	14	15
X	0.32008	0.12801	-0.22099	-0.11673	0.15746
D	-0.12094	-0.17286	-0.12805	0.27294	0.16169
A	0.19543	-0.01042	0.01472	0.10009	0.25320

TABLE 9.27 -- continued

Factor	11	12	13	14	15
RISK	-0.03696	0.15517	-0.05328	0.36381	-0.06798
GSS	0.31576	-0.24533	0.01240	0.05951	-0.06948
OMS	-0.22520	0.21461	0.10693	-0.04652	0.14814
M	0.15822	-0.05462	0.14185	0.02518	0.12264
PQ	-0.29957	0.03046	0.21993	-0.09147	0.15164
L	0.06466	0.35400	-0.22578	-0.15875	0.01682
OPC	-0.00639	0.15802	0.28350	0.03033	-0.09980
BP	0.18946	0.41750	0.14451	0.18472	-0.00217
S	0.00000	-0.00000	0.00000	-0.00000	0.00000
BO	0.31328	-0.00269	0.27097	-0.03495	0.03473
TP	-0.05316	0.16795	-0.14856	0.10896	-0.06195
B	-0.12859	0.06230	-0.07108	-0.04760	0.03615
SP	0.10652	0.05874	0.01048	-0.09678	-0.11597

Notes: Final Communality Estimates total 15.0 and are as follows: 0.0 for S; 1.0 for the rest of the variables.

TABLE 9.28: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 4 - Varimax Rotated Factor Pattern

FACTOR	1	2	3	4	5
X	0.03602	0.04669	0.14883	0.05089	0.00711
D	0.05774	0.09209	0.12082	0.09927	-0.08622
A	-0.10231	-0.24787	-0.14726	0.08419	0.91011
RISK	0.14759	-0.17610	0.06520	0.92643	0.07767
GSS	0.08372	0.00725	0.04552	0.01376	-0.03982
OMS	-0.16731	-0.06523	-0.08774	0.03698	-0.04746
M	-0.02422	-0.01799	-0.10685	0.02106	-0.08576
PQ	0.95101	-0.00952	0.10464	0.13279	-0.08740
L	0.00879	0.05939	-0.16568	0.07772	0.02935
OPC	0.04738	0.10469	0.13445	-0.01761	0.05144
BP	-0.03349	0.08663	0.04612	-0.21501	-0.13151
S	0.00000	0.00000	0.00000	0.00000	0.00000
BO	-0.06461	0.03125	-0.02082	0.04200	-0.04468
TP	0.05751	-0.07338	0.00566	-0.05950	0.03865
B	-0.01518	0.05088	0.03210	-0.02021	-0.11377
SP	0.07991	0.08161	0.04361	0.06382	0.03603

Notes: Final Communality Estimates total 15.0 and are as follows: 0.0 for S; 1.0 for the rest of the variables.

TABLE 9.29: Frequency (as Percentage) of Values of Compensation Variables in the Final Knowledge Base in Experiment 5

COMPENSATION VARIABLE	VALUES OF THE VARIABLE				
	1	2	3	4	5
BASIC PAY	5.9	15.6	18.0	9.7	50.7
SHARE	96.5	1.5	0.9	0.8	0.4
BONUS	43.1	10.8	27.0	11.9	7.1
TERMINAL PAY	89.8	1.9	3.8	2.4	2.0
BENEFITS	70.7	18.1	3.6	3.0	4.6
STOCK	80.6	9.9	4.7	2.5	2.3

TABLE 9.30: Range, Mean and Standard Deviation of Values of Compensation Variables in the Final Knowledge Base in Experiment 5

Variable	Minimum	Maximum	Mean	S.D.
BP	1.0000	5.0000	3.8376590	1.3484571
S	1.0000	5.0000	1.0692112	0.4127470
BO	1.0000	5.0000	2.2910941	1.3171851
TP	1.0000	5.0000	1.2498728	0.8092869
B	1.0000	5.0000	1.5251908	1.0241491
SP	1.0000	5.0000	1.3603053	0.8659971

TABLE 9.31: Correlation Analysis of Values of Compensation Variables in the Final Knowledge Base in Experiment 5

(Spearman Correlation Coefficients in the first row for each variable, Prob > |R| under Ho: Rho=0 in the second)

	BP	S	BO	TP	B	SP
BP	1.00000	-0.05978	0.15907	-0.05684	-0.00131	0.03886
	0.00000	0.0080	0.0001	0.0117	0.9538	0.0850
S	-0.05978	1.00000	-0.00454	0.01208	0.05571	-0.02912
	0.0080	0.00000	0.8408	0.5925	0.0135	0.1970
BO	0.15907	-0.00454	1.00000	0.02932	-0.02295	0.05081
	0.0001	0.8408	0.00000	0.1939	0.3093	0.0243
TP	-0.05684	0.01208	0.02932	1.00000	-0.00354	0.00990
	0.0117	0.5925	0.1939	0.00000	0.8755	0.6611
B	-0.00131	0.05571	-0.02295	-0.00354	1.00000	0.06052
	0.9538	0.0135	0.3093	0.8755	0.00000	0.0073
SP	0.03886	-0.02912	0.05081	0.00990	0.06052	1.00000
	0.0850	0.1970	0.0243	0.6611	0.0073	0.00000

TABLE 9.32: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 5 - Eigenvalues of the Correlation Matrix
Total = 6 Average = 0.375

Factor	1	2	3	4	5	6
Eigenvalue	1.175433	1.073561	1.020839	0.975691	0.924350	0.830127
Difference	0.101872	0.052722	0.045148	0.051341	0.094223	0.830127
Proportion	0.1959	0.1789	0.1701	0.1626	0.1541	0.1384
Cumulative	0.1959	0.3748	0.5450	0.7076	0.8616	1.0000

TABLE 9.33: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 5 - Factor Pattern

FACTOR	1	2	3	4	5	6
X	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
D	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
A	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
RISK	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
GSS	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
OMS	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
M	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
PQ	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
L	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
OPC	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
BP	0.73332	0.17248	0.01164	-0.17230	0.08274	0.62915
S	-0.34869	0.55123	0.04419	-0.37155	0.65919	0.00508
BO	0.56837	0.47214	-0.34889	-0.05745	-0.10818	-0.56330
TP	-0.26373	0.32667	-0.63371	0.58018	0.00086	0.29247
B	-0.27872	0.59673	0.33776	-0.14045	-0.64230	0.14099
SP	0.21404	0.23289	0.61754	0.66957	0.24233	-0.10746

Notes: Final Communality Estimates total 6.0 and are as follows: 1.0 for BP, S, BO, TP, B, and SP; 0.0 for the rest of the variables.

TABLE 9.34: Factor Analysis (Principal Components Method) of the Final Knowledge Base of Experiment 5 - Varimax Rotated Factor Pattern

FACTOR	1	2	3	4	5	6
X	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
D	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
A	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
RISK	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
GSS	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
OMS	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
M	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
PQ	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
L	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
OPC	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
BP	0.07084	-0.03061	0.02061	-0.01879	-0.01860	0.99645
S	-0.00257	0.01489	-0.00424	0.03512	0.99909	-0.01845
BO	0.99737	0.01480	0.00534	0.00245	-0.00259	0.07065
TP	0.01471	0.99928	-0.00697	0.00628	0.01488	-0.03035
B	0.00243	0.00629	0.01183	0.99912	0.03512	-0.01864
SP	0.00531	-0.00697	0.99967	0.01181	-0.00423	0.02042

Notes: Final Communality Estimates total 6.0 and are as follows: 1.0 for BP, S, BO, TP, B, and SP; 0.0 for the rest of the variables.

TABLE 9.35: Summary of Factor Analytic Results for the Five Experiments

Experiment	Number of Significant Factors (Eigenvalue > 1)	Percentage of Total Variation	Total Factors	Average Eigenvalue
1	6	65.35	11	0.6875
2	6	57.47	13	0.8125
3	5	72.54	10	0.6250
4	7	70.78	15	0.9375
5	3	54.50	6	0.3750

TABLE 9.36: Expected Factor Identification of Compensation Variables for the Five Experiments Derived from the Direct Factor Analytic Solution

COMPENSATION VARIABLE	EXPECTED FACTOR IDENTIFICATION				
	EXP 1	EXP 2	EXP 3	EXP 4	EXP 5
BASIC PAY	0.2675	0.1652	0.3053	0.2394	0.3043
SHARE	0.2350	0.2267	0.3081	0.0000	0.3371
BONUS	0.2767	0.2190	0.2325	0.2279	0.3591
TERMINAL PAY	0.2587	0.2467	0.2480	0.2400	0.3529
BENEFITS	0.2619	0.2506	0.2784	0.2104	0.3601
STOCK	0.2757	0.2385	0.2863	0.2054	0.3497

TABLE 9.37: Expected Factor Identification of Compensation Variables for the Five Experiments Derived from the Varimax Rotated Factor Analytic Solution

COMPENSATION VARIABLE	EXPECTED FACTOR IDENTIFICATION				
	EXP 1	EXP 2	EXP3	EXP 4	EXP 5
BASIC PAY	0.1212	0.0795	0.1915	0.1677	0.1667
SHARE	0.1206	0.0915	0.1177	0.0000	0.1661
BONUS	0.1017	0.0881	0.0772	0.0997	0.2095
TERMINAL PAY	0.1122	0.1032	0.1096	0.1234	0.1904
BENEFITS	0.1426	0.0907	0.1511	0.1835	0.1741
STOCK	0.1531	0.0959	0.1109	0.1272	0.1777

TABLE 9.38: Expected Factor Identification of Behavioral and Risk Variables for the Five Experiments Derived from the Direct Factor Pattern

VARIABLE	EXPECTED FACTOR IDENTIFICATION				
	Exp 1	Exp 2	Exp 3	Exp 4	Exp 5
<i>Risk</i>	0.2594	0.2238	0.2490	0.2430	0.0000
Experience	0.2808	0.2518	0.2875	0.2688	0.0000
Education	0.0000	0.0000	0.0000	0.2454	0.0000
Age	0.0000	0.0000	0.0000	0.2275	0.0000
General Social Skills	0.2670	0.2117	0.2688	0.2441	0.0000
Managerial Skills	0.0000	0.2244	0.2757	0.2656	0.0000
Motivation	0.2622	0.2160	0.0000	0.1970	0.0000
Physical Qualities	0.2509	0.2667	0.0000	0.2262	0.0000
Communication Ability	0.0000	0.2489	0.0000	0.2294	0.0000
Other Qualities	0.0000	0.0000	0.0000	0.2396	0.0000

TABLE 9.39: Expected Factor Identification of Behavioral and Risk Variables for the Five Experiments Derived from Varimax Rotated Factor Analytic Solution

VARIABLE	EXPECTED FACTOR IDENTIFICATION				
	Exp 1	Exp 2	Exp 3	Exp 4	Exp 5
<i>Risk</i>	0.1226	0.1153	0.1919	0.0472	0.0000
Experience	0.1015	0.1300	0.2416	0.0337	0.0000
Education	0.0000	0.0000	0.0000	0.0502	0.0000
Age	0.0000	0.0000	0.0000	0.0549	0.0000
General Social Skills	0.1251	0.1016	0.1648	0.0204	0.0000
Managerial Skills	0.0000	0.1030	0.2086	0.0507	0.0000
Motivation	0.1014	0.1453	0.0000	0.0272	0.0000
Physical Qualities	0.0895	0.1260	0.0000	0.1764	0.0000
Communication Ability	0.0000	0.0900	0.0000	0.0374	0.0000
Other Qualities	0.0000	0.0000	0.0000	0.0413	0.0000

CHAPTER 10

REALISTIC AGENCY MODELS

In this chapter we describe Models 4 through 7. These models incorporate realism to a greater extent than previous models. The simulation design of these four models is the same. Each model has 200 simulations conducted with a common set of simulation parameters. The two control variables for the simulations are the number of learning periods and the number of contract renegotiation periods. The learning periods run from 5 to 200, while the contract renegotiation periods run from 5 to 25.

In each learning period, there are a number of contract renegotiation periods. The principal utilizes these periods to collect data about the performance of the agents and the usefulness of her knowledge. In each contract renegotiation period, all the agents are offered new compensation by the principal, which they are at liberty to accept or reject. At the end of a prespecified number of contract renegotiation periods (this number being a control variable), the principal uses the data to initiate the learning process. The learning paradigms are two: the genetic algorithm used in the previous studies, and the specialization-generalization learning operator (described in Sec. 10.2 below). This learning process actually uses the data collected in the contract renegotiation periods to change the principal's knowledge base and bring it in line with

changing conditions (such as the number, characteristics and risk aversion of the agents, and also the nature of exogenous risk).

The timing of the agency problem is as follows:

1. The principal offers a compensation scheme to an agent chosen at random. The principal selects this scheme from out of a large number of possible ones based on her current knowledge base and on her estimate of the agent's characteristics.
2. The agent either accepts or rejects the contract according to his reservation welfare. If the agent rejects the contract, nothing further is done. If all the agents reject the contracts they are offered, the principal seeks a fixed number (here, 5) of new agents. This process continues until an agent accepts a contract.
3. If an agent accepts the contract, he selects an effort level based on his characteristics, the contract, and his private information.
4. Nature acts to render an exogenous environment level.
5. Output occurs as a function of the agent's effort level and the exogenous risk.
6. Sharing of the output between the agent and the principal occurs.
7. The principal reviews the agent's performance and using certain criteria either fires him or continues to deal with him in subsequent periods.

The following are the main features common to these models:

1. The agency is multi-period - a number of periods are simulated.
2. The models are all multi-agent models.
3. The agency is dynamic - agents are hired and fired all the time.

4. Agents not only have the option of leaving during contract negotiation time, but they may also be fired by the principal for poor performance. The performance of the agents is evaluated only at the end of the first contract renegotiation period after every learning episode.

Furthermore, all the models follow the basic LEN model of Spremann. The features of the LEN model are:

1. The principal is risk neutral. So, her utility function is linear (L).
2. The agents are all risk averse. In particular, their utility functions are exponential (E).
3. The exogenous risk is distributed normally around a mean of zero (N).
4. The agents' effort is in $[0, 0.5]$.
5. The agents' disutility of effort is the square root of effort.
6. The output is the sum of the agent's effort and the exogenous risk.
7. The agent's payoff is the compensation less the disutility of effort.
8. The principal's payoff is the output less the compensation.
9. The total output of the agency is separable and is the sum of the outputs derived from the actions of the individual agents.
10. Each agent's output is determined by a random value of the exogenous risk, which may or may not be the same as those faced by the other agents.

Features 1 to 8 are explicit in the original LEN model, while features 9 and 10 are implicit.

The informational characteristics of the models are the same, and are as follows:

1. The agent knows his own characteristics and has access to his private information, both of which affect effort selection. This information is personal to the agent and not shared with the other agents or with the principal.
2. The principal possesses a personal knowledge base which consists of if-then rules. These rules help the principal select compensation schemes based on her estimate of the agent's characteristics. The principal also has available to her an estimate of the agents' characteristics. Some of these estimates are exact (eg. age), while others are close (such estimates are based on some deviation around the true characteristics).
3. The principal can only observe the exogenous risk parameter ex-post. The principal evaluates the performance of each agent in the light of the observed ex-post risk and may decide to fire or retain him.
4. All the agents share a common probability distribution from which their reservation welfare is derived. This distribution (called the *rw-pdf*) is for a random variable which is a sum of the values of the elements of compensation. In Models 4 and 5, this sum ranges from 2 to 10 (since they have two elements of compensation, and each element has 5 possible values from 1 to 5). The *rw-pdf* has a peak value at 4 with probability mass 0.6. In Models 6 and 7, this sum ranges from 6 to 30 (since they have six elements of compensation), and the *rw-pdf* has a probability mass of 0.6 at its peak value of 12. For all the models, the *rw-pdf* is monotonically increasing for values below the peak value, and is monotonically decreasing for values after the peak value.

The experimental design is 2×2 . The first design variable is the number of elements of compensation, and it takes two values - 2 and 6. The second design variable is the policy of evaluation of the agents by the principal. This policy plays a role in firing an agent for lack of adequate performance. One policy evaluates the performance of an agent relative to the other agents. By the use of this policy, an agent is not penalized if his performance is inadequate while the performance of the rest of the agents is also inadequate. In this sense, it is a non-discriminatory policy. The other policy evaluates the performance of an agent without taking into consideration the performance of the other agents. By the use of this policy, an agent is fired if his performance is inadequate with respect to an absolute standard set by the principal, without regard to how the other agents have performed. In this sense, it is a discriminatory policy.

By a discriminatory policy is meant an individualized performance appraisal and firing policy, while a non-discriminatory policy means a relative performance appraisal and firing policy. The words "discriminatory" and "non-discriminatory" will be used in the following discussion only in the sense defined above.

Models 4 and 5 follow the basic LEN model which has only two elements of compensation (basic pay and share of output). In Model 4, the principal evaluates the performance of each agent relative to the performance of the other agents, and hence follows a non-discriminatory firing policy. In Model 5 the principal keeps track of the output she receives from each agent and evaluates each agent on an absolute standard. Hence, she follows a discriminatory policy. Models 6 and 7 follow the basic LEN model, but incorporate four additional elements of compensation (bonus payments,

terminal pay, benefits, and stock participation), as in the previous studies. In Model 6, the principal follows a non-discriminatory evaluation and firing policy, while in Model 7, she follows a discriminatory policy.

The two basic control variables for the simulation are the number of learning periods and the number of contract renegotiation (or data gathering) periods. A number of statistics are collected in these studies, and they are grouped by their ability to address some fundamental questions:

1. The first group of statistics pertains to the simulation methodology. They report the state of the principal's knowledge base. These statistics cover the average and maximum fitness of the rules, their variance around the mean, and the entropy of the normalized fitnesses.
2. The second group of statistics describes the type of compensation schemes offered to the agents by the principal throughout the life of the agency. They report the mean and variance of each element of compensation.
3. The third group of statistics describes the composition of compensation schemes in the final knowledge base of the principal (i.e., at the termination of the simulation). They report the mean and variance of each element of compensation. These statistics differ from those in group two in that they characterize the state of the principal's knowledge base, while those in the second group also capture the compensations activated as a result of the characteristics of the agents who participate in the agency.

4. The fourth group of statistics deals with the movement of agents. They describe the mean and variance of the agents who resigned from the agency on their own (because of the failure of the principal to meet their reservation welfares), those who were fired by the principal (because of their inadequate performance), and those who remained active in the agency when the simulation terminated.
5. The fifth group of statistics deals with agent factors, which measure the change in satisfaction of the agents as they participate in the agency. They help answer the question, "Is the agent better off by participating in this particular model of agency?". These statistics cover the three types of agents - those who resigned, those who were fired, and the remaining agents who are employed (called the "normal" agents).
6. The sixth group of statistics deals with the mean and variance of the satisfaction of the agents, again distinguishing between resigned, fired and normal agents. An agent's satisfaction is calculated from the utility of net income of the agent. However, the term "satisfaction" is used instead of "utility" since the former may take into consideration some intrinsic satisfaction levels which are measured subjectively; see, for example, Chapter 7 on Motivation Theory.
7. The seventh group of statistics reports the mean and variance of the satisfaction level of the agents at termination. For the agents who have resigned, this is the satisfaction derived from an agency period just prior to

resignation. For the agents who have been fired, this is the satisfaction they derived in the agency period they were fired. For normal agents, this is the satisfaction they obtained at the termination of the simulation.

8. The eighth group of statistics covers the mean and variance of the number of agency interactions, reporting separately for resigned, fired and normal agents.
9. The ninth group of statistics details the mean and variance of the number of rules that were activated for each of the three types of agents in the principal's knowledge base.
10. The tenth group of statistics describes the mean and variance of the number of rules that were activated during the final iteration of the simulation.
11. The eleventh group of statistics deals with the principal, and report on the mean and variance of the principal's satisfaction, the principal's factor (which helps answer the question, "Is the principal better off in this agency model?"), and the satisfaction derived by the principal at termination.
12. The twelfth group of statistics details the mean and variance of payoff received by the principal from each of the three kinds of agents. This group of statistics are relevant only in Models 5 and 7, since this information is used by the principal to engage in discriminatory evaluation of the agents' performance.
13. The final group of statistics computes the fit of the principal's knowledge base with the dynamic agency environment. This fit is characterized by a

least squares computation between the antecedents of the principal's knowledge base and the agents' true characteristics, and also between the antecedents and the principal's estimate of the agents' true characteristics.

Whenever some statistic distinguishes between the three types of agents, a statistic that includes all the agents is also computed. These statistics enable one to study the behavior and performance of the agency along several parameters. These statistics are used to study the appropriate correlations.

10.1 Characteristics of Agents

For the purpose of the simulation, the characteristics of the agents are generated randomly from probability distributions. These distributions capture the composition of the agent pool. Other distributions may be used in another agency context. For these studies, some of the characteristics are generated independently of others, while some are generated from conditional distributions. Education, experience and general social skills are conditional on the age of the agent, while office and managerial skills are conditional on the education of the agent. Each agent is an "object" consisting of the following:

1. Nine behavioral characteristics.
2. Three elements of private information.
3. Index of risk aversion, generated randomly from the uniform (0,1) distribution.
4. A vector which plays a role in effort selection by the agent.

The probability distributions are detailed in Table 10.68. The index of risk aversion is unique to the agent and is drawn from the uniform (0,1) distribution. When the principal offers a compensation scheme, the agent draws a reservation welfare from the associated distribution, and compares the utility of the reservation compensation with the utility of the compensation offered by the principal (in these models, the agent does not take into account the expected utility from future contracts). The agent rejects the contract if the latter utility does not exceed the former.

10.2 Learning with Specialization and Generalization.

The structure of the antecedents of the principal's knowledge base have been modified. In the previous models, each antecedent was a single number between 1 and 5 inclusive. However, it is felt that more realism is captured (and the application of other learning operators is made possible) if each antecedent is expressed as an interval bounded inclusively by 1 and 5. This would enable the principal's knowledge base to be as precise or as general as necessary. For example, if the agents who participated in the agency in some learning period had a wide diversity of characteristics, then the knowledge base would be appropriately generalized so that the principal would be able to offer contracts to as many of them as possible. Similarly, if the agents had characteristics which were close to the others, then the principal's knowledge base could be specialized or made more precise in order to distinguish between the agents and tailor compensation schemes appropriately.

This tuning and adjustment is made possible by noting the number of times each antecedent was applicable in the process of searching for an appropriate compensation scheme for each agent. If, during the learning episode, the count of any antecedent of a rule exceeded some average of the counts in the knowledge base, that would imply that that antecedent is too general and is applicable to too many agents. In such a case, the antecedent's length of interval is reduced in a systematic manner. Again, if an antecedent in a rule had a very low count, that would imply that the antecedent of that rule was not being used much. In such a case, the length of interval of that antecedent would be expanded.

The process of reducing the length of an antecedent's interval is called specialization, and the process of increasing it is called generalization. A fixed step size may be associated with each process. We choose the same step size of 0.25 for both the processes. If l_i is the lower bound of an antecedent and u_i is its upper bound, then $(u_i - l_i)$ times the step size (0.25) is the size of the specialization and generalization step size (S). For the above antecedent, the learning operators would act as follows:

Specialization: $l_i \leftarrow l_i + S$; $u_i \leftarrow u_i - S$.

Generalization: $l_i \leftarrow l_i - S$; $u_i \leftarrow u_i + S$.

The step size S is therefore proportional to the length of the interval. Updating of the bounds of the antecedent takes place in each learning episode after the application of the mating operator and before the application of the mutation operator of the genetic algorithm.

10.3 Notation and Conventions

We use the following notation to describe the results for all the models:

The prefix $E[]$ denotes the Mean (expected value), and the prefix $SD[]$ denotes the Standard Deviation;

BP: Basic Pay;

SH: Share;

BO: Bonus;

TP: Terminal Pay;

BE: Benefits;

SP: Stock Participation;

LP: Learning Periods;

CP: Contract Periods;

MAXFIT: Maximum Fitness of Rules;

AVEFIT: Average Fitness of Rules;

VARFIT: Variance of Fitness of Rules;

ENTROPY: Shannon Entropy of Normalized Fitnesses of Rules;

COMP: Total Compensation Package;

FIRED: Agents who were Fired;

QUIT: Agents who Resigned;

NORMAL: Agents who remained until the end;

ALL: All the agents;

Several correlations among the dependent variables, in addition to the thirteen groups listed above, will be presented for each model. The correlations are indicated as "+" for positive correlations and as "-" for negative correlations. All correlations are at the 0.1 significance level. After all the models are summarized, they will be compared and the implications will be discussed. Tables 10.1 through 10.16 cover Model 4. Tables 10.17 through 10.34 cover Model 5, Tables 10.35 through 10.48 cover Model 6, and Tables 10.49 through 10.66 cover Model 7. Table 10.67 compares the four models. Sections 10.4 through 10.7 discuss the results for the four models.

10.4 Model 4: Discussion of Results

Model 4 has two elements of compensation, and the principal does not practice discrimination in evaluating the agents. Increasing the number of learning periods tends to result in a lower contract (by individual contract element and also by total contract) offered to agents. Increasing the number of data collection periods during which no evaluation of the agents occurs, results in a higher contract being offered to the agents. In both cases, the variance of the total contract increases (Table 10.2). Interestingly enough, in both cases, the contracts that make up the final knowledge base of the principal show positive correlation (Table 10.3). Tables 10.2 and 10.3 together imply that while the final knowledge base favors comparatively high contracts, the principal is able to select only the low contracts. This affects the agents' factors (which determine if the agents are better off at termination in this agency model) adversely. In fact, increasing the number of learning periods leaves all the agents worse off (Table 10.5).

However, this adverse affect is not carried through to the agents' satisfactions. All the agents seem more satisfied the longer the agency process. But the agents' satisfactions decreased on average by increasing the number of data collection periods (Table 10.6). In other words, observability by the principal affects their satisfaction adversely. This may be due to the fact that the principal has more data with which she can measure the usefulness of her knowledge base (by measuring the relative importance of each of the antecedent clauses in the rules), thus allowing her to tailor compensation schemes (which form the consequent clauses in the rules of her knowledge base) to reward agents accurately. Because of the fundamentally adversarial relationship between the agents and the principal, this would decrease the mean satisfaction of the agents.

The mean agent factor for fired agents is positively correlated with the mean satisfaction of the principal (Table 10.13). However, the agents' satisfaction and the principal's satisfaction held inverse relationships, except in the case of normal agents (Table 10.14). This is consistent with the fact that the agents who quit and those who were fired obtained, on average, less satisfaction than the normal agents. However, because of the extremely dynamic environment (a mean of 444 agents across all the simulations), the overall mean satisfaction of the agents is negatively correlated with the principal's satisfaction (Table 10.67).

With increase in the length of the simulation, more agents quit and were fired. However, the expected number of agents fired decreased. In all, a mean of 5 agents were fired (Table 10.67).

10.5 Model 5: Discussion of Results

Model 5 has two elements of compensation, and the principal evaluates the performance of agents in a discriminatory manner. The value of individual elements of the contract as well as the value of the total contract offered to agents decreased with increases in the number of learning periods. When the number of contract periods for each learning period was increased, only the mean share offered to the agents increased, but no significant results were available for the rest of the elements of contract. The variance of the total contract increased both times (Table 10.18). The principal's final knowledge base is also consistent with this result (Table 10.19).

Increasing the number of learning periods left the agents worse off at termination, while increasing the number of contract periods merely decreased the variance of the agent factors (Table 10.21). However, the agents' satisfaction showed positive correlation with the number of learning periods, except for agents who were fired (Table 10.22). This positive correlation also extends to those agents who quit of their own accord. This implies that while the satisfactions rose with more learning periods, they did not rise high enough or in a timely way for some agents. Again, as in Model 4, increasing observability (number of contract periods) by the principal correlated negatively with agents' satisfactions, while decreasing their variance.

In Model 5, payoff from individual agents is known, which enables the principal to practice discrimination in firing agents. The mean payoff of agents who quit, of those who stayed on (normal agents), and also of all the agents (considered as a whole), showed positive correlation with the number of learning periods, while the number of

contract periods correlated negatively with all but normal agents. For agents who were fired, there were no significant correlations. This implies that on the whole, observability by the principal affects the agents' payoffs adversely (Table 10.30).

The principal's satisfaction also correlated negatively with the mean outcomes from the agents (Table 10.33). Mean payoff from an agent may increase with an increase in the number of learning periods while the outcome from that agent decreases because the principal is offering smaller contracts.

The mean satisfaction of the two parties showed positive correlation only in the case of agents who were fired and normal agents. There is an inverse relationship between the mean satisfaction of the principal and the mean satisfaction of agents who quit. This is also true in the case of all the agents (taken as a whole) (Table 10.31). This implies that while the principal's satisfaction was high, most of the contribution came from agents who ultimately resigned from the agency, while those who were fired used less effort and had commensurately higher contracts. This may suggest the reason for why some agents quit and why some agents were fired. On the whole, this Model is extremely dynamic since the total number of agents who quit (996) and the total number of agents who were fired (16) is the highest for all the four models (Table 10.67).

10.6 Model 6: Discussion of Results

Model 6 has six elements of compensation, and the principal does not practice any discrimination in evaluating the performance of the agents. As with the previous Models

4 and 5, compensation offered to agents correlated negatively with the number of learning periods (Table 10.36). The value of compensation in the final knowledge base of the principal also correlated negatively with both the number of learning periods and the number of contract periods (Table 10.37).

The mean principal's satisfaction correlated negatively with the mean satisfaction of agents who quit and all the agents (considered as a whole). There were no corresponding significant correlations between the principal's satisfaction and the agents' factors (Table 10.47). The principal's factor (which indicates if she is better off by participating in this agency model) and factors of agents who were fired correlate negatively. This of course explains why these agents were fired. The agency environment is more stable than the previous two models. Only 234 agents quit, while only 4 agents were fired (Table 10.67).

10.7 Model 7: Discussion of Results

Model 7 has six elements of compensation, and the principal practices discrimination in her evaluation of the agents. As in the previous models, a higher number of learning periods is associated with lower value of compensation packages offered to the agents (Table 10.50). A higher number of contract periods correlates negatively with mean basic pay, but positively with mean value of stock participation (no significant correlations were observed at the 0.1 level for the other elements of compensation) (Table 10.50). In the final knowledge base of the principal, the variances of the elements of compensation showed negative correlation with the number of learning

periods, while the mean values for some of the elements of compensation (share of output and stock participation) and the total contract correlated positively with the number of contract periods (Table 10.51).

The principal has available to her (in this Model 7) the payoffs from each agent. The mean payoff from agents who quit and all the agents taken as a whole showed positive correlation with the number of learning periods, while there was negative correlation for the mean payoff from fired agents. This implies that the principal succeeded in learning to control effort selection by the agents in such a way as to increase the payoff. This need not imply that the agents are better off, or that their mean satisfaction is high (see discussion in the next paragraph). The number of contract periods correlated negatively with the mean payoff from all types of agents except for fired agents (who had no significant correlation) (Table 10.59). This may seem counter-intuitive, since having more data should lead to better control. However, collecting data takes time. The longer it takes time, the longer the principal defers using the learning mechanisms. This gives the agents time to get away with a smaller contribution to payoff, while collecting commensurately larger contracts until the principal learns.

The mean agent factor for fired agents, and the mean satisfaction of agents who quit and of normal agents, correlates negatively with the mean satisfaction of the principal (Tables 10.62 and 10.63). The principal is also able to observe the outcomes from each agent individually. The mean outcome from all the types of agents (except those who were fired, for whom there are no significant correlations at the 0.1 level)

correlates negatively with the principal's satisfaction (Table 10.65). Observability, in this case also, is detrimental to the interests of the agents.

10.8 Comparison of the Models

Table 10.67 summarizes the key statistics of the four models. The contracts offered to the agents by the principal are higher in value in Models 6 and 7 (than in Models 4 and 5) where the number of elements are more (six, as compared to two in Models 4 and 5). However, the value of the contract per element of compensation (i.e. the normalized statistic) is the highest in Model 4 (two elements of compensation and non-discriminatory evaluation), followed by Models 6 and 7 (Table 10.67). This suggests that in the absence of complex contracts and individualized observability, the principal can only offer higher contracts in an effort to second-guess the reservation welfare of the agents and to retain the services of good agents. Again, since observability is poor, the principal can only offer comparatively higher contracts to all the agents. Increasing either the complexity of contracts or the observability enables the principal to be more efficient. However, the principal must have an instrument capable of being flexible in order to be efficient. This is not possible when the contracts are very simple, even if the principal is able to observe each agent individually. Hence, in Model 5, the principal can only effectively punish poor performance. If she attempts to reward good performance using only two elements of compensation (basic pay and share of output), her own welfare is affected. Hence, the value of contracts in Model 5 is uniformly lower than in the other models. This also leads us to expect that the

reservation welfare of many agents may not be met. Table 10.67 confirms this. The number of agents who quit of their own accord is the highest of all models. Similarly, the principal is unable to induce proper effort selection using only two elements of compensation. However, this does not stop her from punishing (effectively and individually) poor performers. This leads us to expect that the number of agents fired in Model 5 would be highest of all the models. Table 10.67 again confirms this expectation.

Agent factors indicate whether the agents were better off or worse off on the whole in the particular agency model (with positive factors indicating better off and negative factors indicating worse off). This is a measure of the difference in satisfaction enjoyed by the agents normalized for number of learning periods and contract periods. Agents were better off to a greater extent when the number of compensation elements were two rather than six, and when the principal practiced non-discriminatory evaluation of agents' performance. This is because the principal has less scope for controlling agents' effort selection through complex contracts, and no individualized evaluation of agents' performances and hence no possibility of penalizing agents with poor performance. Therefore, in all cases except Model 7, agents as a whole were better off.

Looking at specific types of agents, the agents who quit were better off in the non-discriminatory cases (Models 4 and 6), and in the case of two elements of compensation (Models 4 and 5). The same holds true for agents who were fired. However, for normal agents, the greatest increase in satisfaction (compared across the models) occurred in Model 5 (two elements of compensation with discriminatory

evaluation). This implies that the principal discriminated in favor of the normal agents, while terminating the services of undesirable agents and forcing other agents to quit by offering very low contracts. Normal agents suffered the most in Model 4 where the number of elements of compensation are two, and the principal does not practice discrimination. This means that the complexity of compensation plans is insufficient to selectively reward good agents in Models 4 and 5, while a non-discriminatory evaluation practice (as in Models 4 and 6) adds to the unfavorable atmosphere for the good agents. It appears that increasing either the number of elements of compensation or practicing a discriminatory evaluation is sufficient to selectively reward good agents (as in Models 5 and 6), but the dual approach does not help the normal agents (as in Model 7), even though their factors are almost double when compared to Model 4.

Interestingly, the greatest increase in satisfaction is observed for the agents who were fired for all the models. At the same time, their mean satisfaction is the lowest. This implies that their payoff contribution to the principal is also low. In the case of two elements of compensation and non-discrimination (Models 4 and 5), the agents who were eventually fired took advantage of the principal's inability to focus on their performance, thereby increasing their satisfaction at a rate which was higher than that of other types of agents. Whenever the principal had complex contracts as a manipulating tool (as in Models 6 and 7), or whenever she had sufficient information to evaluate performances discriminatively (as in Models 5 and 7), the factors for fired agents show a significant decline. This implies that they were fired sooner before they could increase their own payoffs to the extent they could in the other models.

The mean satisfaction of agents showed significant increase (about 70%) in Models 6 and 7 over Models 4 and 5. Comparing with the drop in agent factors in Models 6 and 7 over Models 4 and 5, this implies that using more elements of compensation raises the level of satisfaction by about 70%, but does not cause a comparatively higher rise in satisfaction as the agency progresses. When the number of elements of compensation are two (Models 4 and 5), the mean satisfaction of agents is higher in the discriminatory case compared to the non-discriminatory case, except (of course) for agents who were eventually fired. However, when the number of elements of compensation were increased to six (Models 6 and 7), all agents experienced decreased mean satisfaction in the discrimination case (Model 7). This seems to suggest that complexity of contracts and the practice of discrimination work at cross purposes in satisfying all agents.

On the one hand, if the goal of the agency is to rapidly improve satisfaction levels (or increase the rate of their improvement), then discrimination is the best policy (since Model 5 has the highest agent factors if the factors for fired agents is ignored). Such a goal might be reasonable for an existing agency currently suffering from low satisfaction levels or low profit levels. A discrimination policy would get rid of shirking agents, convey a motivational message to good agents, and increase profits by paring down the value of contracts temporarily.

On the other hand, if the goal of the agency is to achieve a high mean satisfaction level, attract better agents by matching the general reservation welfare, and decrease agent turnover, then a non-discriminatory evaluation policy coupled with complex

contracts is the best policy (as in Model 6). Such a model would be very useful if the cost of human resource management is a significant expense item for the principal. Further, while a high satisfaction level is achieved in Model 6, further increase in satisfaction is only gradual. The emphasis is not on agent factors. In many real-world situations, if the initial satisfaction level of agents is high, further attempts to increase that level might yield diminishing returns.

In Table 10.67, negative values for the mean satisfaction of agents occur because the satisfaction of the agents is dependent on their risk-aversion of net income, which is modeled as a negative exponential function which always takes negative values. Hence, the absolute values in any model taken by themselves do not convey much information. They must be compared with the values from the other models.

The mean satisfaction of the principal is greatest in Model 6 (six elements of compensation with non-discriminatory evaluation), and least in Model 5 (two elements of compensation with discriminatory evaluation). Her mean satisfaction is higher the more complex the contracts (since this allows her to tailor compensation to a wide variety of agents), and it is also higher if she does not practice discrimination. So, while discrimination is good for some agents in some circumstances, it is never a desirable policy for the principal. When complex contracts are involved, the practice of discrimination erodes the mean satisfaction of all parties only marginally, while decreasing agents' factors significantly.

The greatest improvement in satisfaction for the principal, however, takes place in Model 5 (two elements of compensation and discriminatory policy). The improvement

is marginally positive in Model 7 (complex contracts and discrimination), while it is marginally negative in Model 6 (complex contracts and no discrimination). Hence, depending on the goals of the agency vis-a-vis satisfaction of the principal, Model 5 (which ensures greatest rate of increase of satisfaction) or Model 6 (which ensures highest mean satisfaction) may be chosen.

Predictably, the greatest number of agents were fired in the case of discriminatory evaluation (Models 5 and 7), while the greatest number of agents quit in Model 5, followed by Model 4. This implies that in Model 4, some of the agents were not satisfied with the simple contracts offered to them (which did not meet their reservation levels), while in Model 5, the principal forced some of the poorly performing agents to resign by assigning them comparatively low contracts. The use of complex contracts significantly reduces the number of agents who quit and also the number of agents who were fired. This is because complex contracts enable the principal tailor contracts efficiently to as many agents as possible. This ensures a more stable agency environment.

10.9 Examination of Learning

There is no significant advantage in conducting a longer simulation in order to increase maximum fitness of rules. Only uniformity of rule fitnesses (denoted by entropy) is better achieved through longer simulations. Increasing the length of the contract period increases the maximum fitness while also increasing the variance (Tables 10.1, 10.17, 10.35, and 10.49). Only in the case of Model 5 the average fitness showed

positive correlation with number of learning periods, but not at the 0.1 level of significance. In all other cases, the correlation was negative, but not at the 0.1 level of significance (except for Model 6, which showed significance). This suggests that the models may be GA-deceptive. Further study is necessary to verify this, and suggestions are made in Chapter 12 (Future Research). Another reason for this behavior may be due to the fact that the functions which calculate fitness of rules do not cover all the factors that cause the fitness to change. Of necessity, the agents' private information must remain unknown to the principal's learning mechanism. Further, the index of risk aversion of the agents is uniformly distributed in the interval (0,1). Computation of fitness is hence not only probabilistic, but also "incomplete".

TABLE 10.1: Correlation of LP and CP with Simulation Statistics (Model 4)

	AVEFIT	MAXFIT	VARIANCE	ENTROPY
LP		-	-	+
CP	-	+	+	-

TABLE 10.2: Correlation of LP and CP with Compensation Offered to Agents (Model 4)

	E[BP]	SD[BP]	E[SH]	SD[SH]	E[COMP]	SD[COMP]
LP	-	-	-		-	+
CP			+	+	+	+

TABLE 10.3: Correlation of LP and CP with Compensation in the Principal's Final KB (Model 4)

	E[BP]	SD[BP]	E[SH]	SD[SH]	E[COMP]	SD[COMP]
LP		-	+	-	+	-
CP		+		+		+

TABLE 10.4: Correlation of LP and CP with the Movement of Agents (Model 4)

	QUIT	E[QUIT]	SD[QUIT]	FIRE	E[FIRE]	SD[FIRE]
LP	+	+	+	+	-	-
CP	+	-	+		-	-

TABLE 10.5: Correlation of LP with Agent Factors (Model 4)

	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	SD[ALL]
LP		-	-	+	-

TABLE 10.6: Correlation of LP and CP with Agents' Satisfaction (Model 4)

	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	E[NORMAL]	E[ALL]	SD[ALL]
LP	+		-	+	+	+	
CP	-	+				-	+

TABLE 10.7: Correlation of LP and CP with Agents' Satisfaction at Termination (Model 4)

	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	E[ALL]	SD[ALL]
LP			-	+	+	
CP	-	+			-	+

TABLE 10.8: Correlation of LP and CP with Agency Interactions (Model 4)

	E[QUIT]	SD[QUIT]	SD[FIRE]	E[NORMAL]	E[ALL]	SD[ALL]
LP	-	-	+	-	-	-
CP	+	+	+		+	+

TABLE 10.9: Correlation of LP with Rule Activation (Model 4)

	E[QUIT]	SD[QUIT]	E[FIRE]	E[ALL]	SD[ALL]
LP	-	-	-	-	-

TABLE 10.10: Correlation of LP with Rule Activation in the Final Iteration (Model 4)

	E[QUIT]	SD[QUIT]	E[ALL]	SD[ALL]
LP	-	-	-	-

TABLE 10.11: Correlation of LP and CP with Principal's Satisfaction and Least Squares (Model 4)

	E[SATP]	SD[SATP]	LASTSATP	FACTOR	BEH-LS	EST-LS
LP	-	+	-	-	+	+
CP	+			-	+	+

TABLE 10.12: Correlation of Agent Factors with Agent Satisfaction (Model 4)

AGENT FACTORS	AGENT SATISFACTION			
	SD[QUIT]	SD[FIRE]	SD[NORMAL]	SD[ALL]
SD[QUIT]	+			
SD[FIRE]		+		
SD[NORMAL]			+	
SD[ALL]				+

TABLE 10.13: Correlation of Principal's Satisfaction with Agent Factors (Model 4)

PRINCIPAL'S SATISFACTION	AGENTS' FACTORS				
	SD[QUIT]	E[FIRE]	SD[FIRE]	SD[NORMAL]	SD[ALL]
E[SATISFACTION]	+	+	-	+	+
SD[SATISFACTION]				-	

TABLE 10.14: Correlation of Principal's Satisfaction with Agents' Satisfaction (Model 4)

PRINCIPAL'S SATISFACTION	AGENTS' SATISFACTION				
	E[QUIT]	SD[QUIT]	SD[NORMAL]	E[ALL]	SD[ALL]
E[SATISFACTION]	-		+	-	+
SD[SATISFACTION]	-	+	-		+

TABLE 10.15: Correlation of Principal's Last Satisfaction with Agents' Last Satisfaction (Model 4)

	AGENTS' LAST SATISFACTION		
	SD[FIRE]	E[NORMAL]	SD[NORMAL]
PRINCIPAL'S LAST SATISFACTION	-	-	+

TABLE 10.16: Correlation of Principal's Factor with Agent Factors (Model 4)

	SD[QUIT]	E[FIRE]	SD[NORMAL]
PRINCIPAL'S FACTOR	+	+	+

TABLE 10.17: Correlation of LP and CP with Simulation Statistics (Model 5)

	AVEFIT	MAXFIT	VARIANCE	ENTROPY
LP		-	-	+
CP	-	+	+	-

TABLE 10.18: Correlation of LP and CP with Compensation Offered to Agents (Model 5)

	E[BP]	SD[BP]	E[SH]	SD[SH]	E[COMP]	SD[COMP]
LP	-		-	-	-	+
CP			+	+		+

TABLE 10.19: Correlation of LP and CP with Compensation in the Principal's Final Knowledge Base (Model 5)

	E[BP]	SD[BP]	E[SH]	SD[SH]	E[COMP]	SD[COMP]
LP		-		-		-
CP		+		+		+

TABLE 10.20: Correlation of LP and CP with the Movement of Agents (Model 5)

	QUIT	E[QUIT]	SD[QUIT]	FIRE	E[FIRE]	SD[FIRE]
LP	+	+	+	+	-	-
CP	+	-	+		-	-

TABLE 10.21: Correlation of LP with Agent Factors (Model 5)

	SD[QUIT]	E[FIRE]	SD[NORMAL]	SD[ALL]
LP	-	-	-	-
CP	+			+

TABLE 10.22: Correlation of LP and CP with Agents' Satisfaction (Model 5)

	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	SD[NORMAL]	E[ALL]	SD[ALL]
LP	+	-	-	+	-	+	-
CP	-	+		+		-	+

TABLE 10.23: Correlation of LP and CP with Agents' Satisfaction at Termination (Model 5)

	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	SD[NORMAL]	E[ALL]	SD[ALL]
LP	+	-	-	+	-	+	-
CP	-	+		+	-	-	+

TABLE 10.24: Correlation of LP and CP with Agency Interactions (Model 5)

	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	E[NORMAL]	E[ALL]	SD[ALL]
LP	-	-	-	+	-	-	-
CP	+	+		+		+	+

TABLE 10.25: Correlation of LP with Rule Activation (Model 5)

	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	E[ALL]	SD[ALL]
LP	-	-	-	-	-	-

TABLE 10.26: Correlation of LP with Rule Activation in the Final Iteration (Model 5)

	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	E[ALL]	SD[ALL]
LP	-	-	-	-	-	-

TABLE 10.27: Correlation of LP and CP with Payoffs from Agents (Model 5)

	E[QUIT]	SD[QUIT]	SD[FIRE]	E[NORMAL]	SD[NORMAL]	E[ALL]	SD[ALL]
LP	+	-		+	-	+	-
CP	-	+	+		+	-	+

TABLE 10.28: Correlation of LP and CP with Principal's Satisfaction, Principal's Factor and Least Squares (Model 5)

	E[SATP] ¹	SD[SATP]	LASTSATP ²	FACTOR ³	BEH-LS ⁴	EST-LS ⁵
LP	-		-	-	+	+
CP	+	+		-		

¹ SATP: Principal's Satisfaction² Principal's Satisfaction at Termination³ Principal's Factor⁴ Least Squares Deviation from Agents' True Behavior⁵ Least Squares Deviation from Principal's Estimate of Agents' Behavior

TABLE 10.29: Correlation of Agent Factors with Agent Satisfaction (Model 5)

AGENT FACTORS	AGENT SATISFACTION			
	SD[QUIT]	SD[FIRE]	SD[NORMAL]	SD[ALL]
SD[QUIT]	+			
SD[FIRE]		+		
SD[NORMAL]			+	
SD[ALL]				+

TABLE 10.30: Correlation of Principal's Satisfaction with Agent Factors (Model 5)

PRINCIPAL'S SATISFACTION	AGENTS' FACTORS				
	SD[QUIT]	E[FIRED]	SD[FIRED]	SD[NORMAL]	SD[ALL]
E[SATISFACTION]	+	+	-	+	+
SD[SATISFACTION]	+	-	+		+

TABLE 10.31: Correlation of Principal's Satisfaction with Agents' Satisfaction (Model 5)

PS ¹	AGENTS' SATISFACTION						
	E[QUIT]	SD[QUIT]	E[FIRED]	SD[FIRED]	SD[NORMAL]	E[ALL]	SD[ALL]
²	-	+	+	-	+	-	+
³	-	+				-	+

¹ PS: This column contains the mean and standard deviation of the Principal's Satisfaction

² Mean Principal's Satisfaction

³ Standard Deviation of Principal's Satisfaction

TABLE 10.32: Correlation of Principal's Last Satisfaction with Agents' Last Satisfaction (Model 5)

	AGENTS' LAST SATISFACTION					
	E[QUIT]	E[FIRED]	E[NORMAL]	SD[NORMAL]	E[ALL]	SD[ALL]
PRINCIPAL'S LAST SATISFACTION	-	+	-	+	-	+

TABLE 10.33: Correlation of Principal's Satisfaction with Outcomes from Agents (Model 5)

PS ¹	OUTCOMES FROM AGENTS							
	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	⁴	⁵	E[ALL]	SD[ALL]
²	-	+	-	+	-	+	-	+
³		-		-		-		-

¹ PS: This column contains the mean and standard deviation of the Principal's Satisfaction

² Mean Principal's Satisfaction

³ Standard Deviation of Principal's Satisfaction

⁴ E[NORMAL] : Mean Outcome from Normal (non-terminated) Agents

⁵ SD[NORMAL]

TABLE 10.34: Correlation of Principal's Factor with Agents' Factors (Model 5)

	E[FIRE]	SD[FIRE]	SD[NORMAL]
PRINCIPAL'S FACTOR	+	-	+

TABLE 10.35: Correlation of LP and CP with Simulation Statistics (Model 6)

	AVEFIT	MAXFIT	VARIANCE	ENTROPY
LP	-	-	-	+
CP	-	+	+	-

TABLE 10.36: Correlation of LP and CP with Compensation Offered to Agents (Model 6)

	E ¹	SD ¹	SD ²	SD ³	E ⁴	SD ⁴	SD ⁵	E ⁶	SD ⁶	E ⁷	SD ⁷
LP	-	-	-	-	-	-	-	-	-	-	+
CP	-			+			+				+

¹ BASIC PAY; ² SHARE OF OUTPUT; ³ BONUS PAYMENTS; ⁴ TERMINAL PAY

⁵ BENEFITS; ⁶ STOCK PARTICIPATION; ⁷ TOTAL CONTRACT

TABLE 10.37: Correlation of LP and CP with Compensation in the Principal's Final Knowledge Base (Model 6)

	E ¹	SD ¹	SD ²	SD ³	SD ⁴	SD ⁵	SD ⁶	SD ⁷
LP		-	-	-	-	-	-	-
CP	-	-	-		-		-	-

¹ BASIC PAY; ² SHARE OF OUTPUT; ³ BONUS PAYMENTS; ⁴ TERMINAL PAY

⁵ BENEFITS; ⁶ STOCK PARTICIPATION; ⁷ TOTAL CONTRACT

TABLE 10.38: Correlation of LP and CP with the Movement of Agents (Model 6)

	QUIT	E[QUIT]	SD[QUIT]	FIRE	E[FIRE]	SD[FIRE]
LP	+	+	+	+		
CP	+		+		-	-

TABLE 10.39: Correlation of LP and CP with Agent Factors (Model 6)

	SD[QUIT]	E[NORMAL]	SD[ALL]
LP	-		-
CP		+	

TABLE 10.40: Correlation of LP and CP with Agents' Satisfaction (Model 6)

	SD[QUIT]	SD[FIRE]
LP		+
CP	+	

TABLE 10.41: Correlation of LP and CP with Agents' Satisfaction at Termination (Model 6)

	SD[QUIT]	SD[FIRED]	E[ALL]	SD[ALL]
LP		+	+	
CP	+			+

TABLE 10.42: Correlation of LP and CP with Agency Interactions (Model 6)

	E[QUIT]	SD[QUIT]	E[FIRED]	SD[FIRED]	E[NORMAL]	E[ALL]	SD[ALL]
LP	-	-		+	-	-	-
CP		+	+				+

TABLE 10.43: Correlation of LP and CP with Rule Activation (Model 6)

	E[QUIT]	SD[QUIT]	E[FIRED]	E[ALL]	SD[ALL]
LP	-	-	-	-	-
CP	-				

TABLE 10.44: Correlation of LP and CP with Rule Activation in the Final Iteration (Model 6)

	E[QUIT]	SD[QUIT]	E[FIRED]	SD[FIRED]	E[ALL]	SD[ALL]
LP	-	-			-	-
CP			+	+		

TABLE 10.45: Correlation of LP and CP with Principal's Satisfaction and Least Squares (Model 6)

	E[SATP ¹]	SD[SATP]	LASTSATP ²	BEH-LS ³	EST-LS ⁴
LP	-	+	-		
CP		+		+	+

¹ SATP: Principal's Satisfaction² Principal's Satisfaction at Termination³ Least Squares Deviation from Agents' True Behavior⁴ Least Squares Deviation from Principal's Estimate of Agents' Behavior

TABLE 10.46: Correlation of Agents' Factors with Agents' Satisfaction (Model 6)

1	AGENTS' SATISFACTION						
	2	3	4	5	6	E[ALL]	SD[ALL]
2	+						
4		-					
4			+				
6				-			
6					+		
E[ALL]						-	
SD[ALL]							+

¹ This column denotes Agents' Factors² Standard Deviation of Factor/Satisfaction of Agents who Quit³ Mean Factor/Satisfaction of Agents who were Fired⁴ Standard Deviation of Factor/Satisfaction of Agents who were Fired⁵ Mean Factor/Satisfaction of Agents who remained Active (Normal)⁶ Standard Deviation of Factor/Satisfaction of Agents who remained Active (Normal)

TABLE 10.47: Correlation of Principal's Satisfaction with Agents' Factors and Agents' Satisfaction (Model 6)

PRINCIPAL'S SATISFACTION	AGENTS' FACTORS		AGENTS' SATISFACTION	
	E[QUIT]	SD[NORMAL]	E[QUIT]	E[ALL]
E[SATISFACTION]		+	-	-
SD[SATISFACTION]	+		+	+

TABLE 10.48: Correlation of Principal's Factor with Agents' Factor (Model 6)

	E[FIRE]	SD[FIRE]	SD[NORMAL]
PRINCIPAL'S FACTOR	-	-	+

TABLE 10.49: Correlation of LP and CP with Simulation Statistics (Model 7)

	AVEFIT	MAXFIT	VARIANCE	ENTROPY
LP		-		+
CP	-	+	+	-

TABLE 10.50: Correlation of LP and CP with Compensation Offered to Agents (Model 7)

	E ¹	SD ¹	E ²	SD ²	E ³	SD ³	SD ⁴	SD ⁵	E ⁶	SD ⁶	E ⁷	SD ⁷
LP	-	-	-	-	-	-	-	-		-	-	+
CP	-							+	+	+		+

¹ BASIC PAY; ² SHARE OF OUTPUT; ³ BONUS PAYMENTS; ⁴ TERMINAL PAY;

⁵ BENEFITS; ⁶ STOCK PARTICIPATION; ⁷ TOTAL CONTRACT

TABLE 10.51: Correlation of LP and CP with Compensation in the Principal's Final Knowledge Base (Model 7)

	SD ¹	E ²	SD ²	SD ³	SD ⁴	SD ⁵	E ⁶	SD ⁶	E ⁷	SD ⁷
LP	-		-	-	-	-		-		-
CP		+					+		+	

¹ BASIC PAY; ² SHARE OF OUTPUT; ³ BONUS PAYMENTS; ⁴ TERMINAL PAY;

⁵ BENEFITS; ⁶ STOCK PARTICIPATION; ⁷ TOTAL CONTRACT

TABLE 10.52: Correlation of LP and CP with the Movement of Agents (Model 7)

	QUIT	E[QUIT]	SD[QUIT]	FIRED	E[FIRED]	SD[FIRED]
LP	+	+	+	+		
CP	+	-	+		-	-

TABLE 10.53: Correlation of LP with Agent Factors (Model 7)

	SD[QUIT]	SD[NORMAL]	SD[ALL]
LP	-	-	-

TABLE 10.54: Correlation of LP and CP with Agents' Satisfaction (Model 7)

	E[QUIT]	SD[QUIT]	SD[FIRED]	SD[ALL]
LP	+		+	
CP		+		+

TABLE 10.55: Correlation of LP and CP with Agents' Satisfaction at Termination (Model 7)

	SD[QUIT]	SD[FIRE]	E[ALL]
LP		+	+
CP	+		

TABLE 10.56: Correlation of LP and CP with Agency Interactions (Model 7)

	E[QUIT]	SD[QUIT]	SD[FIRE]	E[NORMAL]	E[ALL]	SD[ALL]
LP	-	-	+	-	-	-
CP	+	+				+

TABLE 10.57: Correlation of LP and CP with Rule Activation (Model 7)

	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	E[ALL]	SD[ALL]
LP	-	-	-	-	-	-
CP	-					

TABLE 10.58: Correlation of LP with Rule Activation in the Final Iteration (Model 7)

	E[QUIT]	SD[QUIT]	SD[FIRE]	E[ALL]	SD[ALL]
LP	-	-	+	-	-

TABLE 10.59: Correlation of LP and CP with Payoffs from Agents (Model 7)

	E[QUIT]	SD[QUIT]	E[FIRED]	SD[FIRED]	E ¹	SD ¹	E[ALL]	SD[ALL]
LP	+	-	-	+		-	+	-
CP	-	+			-		-	+

¹ NORMAL (Active) Agents

TABLE 10.60: Correlation of LP and CP with Principal's Satisfaction (Model 7)

	E[SATP ¹]	SD[SATP]	LASTSATP ²
LP	-	+	-

¹ SATP: Principal's Satisfaction² Principal's Satisfaction at Termination

TABLE 10.61: Correlation of Agent Factors with Agent Satisfaction (Model 7)

AGENTS' FACTOR	AGENTS' SATISFACTION				
	SD[QUIT]	E[FIRED]	SD[FIRED]	SD[NORMAL]	SD[ALL]
SD[QUIT]	+				
E[FIRED]		-			
SD[FIRED]			+		
SD[NORMAL]				+	
SD[ALL]					+

TABLE 10.62: Correlation of Principal's Satisfaction with Agent Factors (Model 7)

PRINCIPAL'S SATISFACTION	AGENTS' FACTORS			
	E[FIRED]	SD[FIRED]	SD[NORMAL]	SD[ALL]
E[SATISFACTION]	-	-	+	+
SD[SATISFACTION]			-	

TABLE 10.63: Correlation of Principal's Satisfaction with Agents' Satisfaction (Model 7)

PRINCIPAL'S SATISFACTION	AGENTS' SATISFACTION		
	E[QUIT]	SD[FIRE]	E[ALL]
E[SATISFACTION]	-	-	-
SD[SATISFACTION]	+	+	+

TABLE 10.64: Correlation of Principal's Last Satisfaction with Agents' Last Satisfaction (Model 7)

	AGENTS' LAST SATISFACTION				
	E[QUIT]	E[NORMAL]	SD[NORMAL]	E[ALL]	SD[ALL]
PRINCIPAL'S LAST SATISFACTION	-	-	+	-	+

TABLE 10.65: Correlation of Principal's Satisfaction with Outcomes from Agents (Model 7)

PS ¹	OUTCOMES FROM AGENTS							
	E[QUIT]	SD[QUIT]	E[FIRE]	SD[FIRE]	⁴	⁵	E[ALL]	SD[ALL]
²	-	+		+	-	+	-	+
³	+	-	+	-	+	-	+	-

¹ PS: This column contains the mean and standard deviation of the Principal's Satisfaction

² Mean Principal's Satisfaction

³ Standard Deviation of Principal's Satisfaction

⁴ E[NORMAL] : Mean Outcome from Normal (non-terminated) Agents

⁵ SD[NORMAL]

TABLE 10.66: Correlation of Principal's Factor with Agents' Factor (Model 7)

	E[FIRE]
PRINCIPAL'S FACTOR	+

TABLE 10.67: Comparison of Models
(Standard Deviation in Parenthesis)

MODEL #	4	5	6	7
DESCRIPTION	Non-Discriminatory	Discriminatory	Non-Discriminatory	Discriminatory
COMPENSATION ELEMENTS	2	2	6	6
VARIABLES	78	86	94	102
SIMULATION STATISTICS				
Average Fitness	10401 (8194)	10342 (7432)	10255 (6637)	9263 (5113)
Maximum Fitness	46089 (15960)	44161 (16331)	36123 (16528)	35330 (16221)
Variance of Fitness	0.9803 (0.000049)	0.9803 (0.000041)	0.9803 (0.000040)	0.9803 (0.000039)
Entropy of Fitness	4.4687 (0.1303)	4.4830 (0.1201)	4.5032 (0.1256)	4.4833 (0.1390)
Contract Offered to Agents	5.7835 (0.8466)	5.3706 (1.0095)	17.0137 (1.8618)	16.9674 (1.1063)
Contract Offered to Agents (Normalized)	2.8918	2.6853	2.8356	2.8279
MOVEMENT OF AGENTS				
Total Agents Who Quit	444	996	234	232
Total Agents Fired	5	16	4	7

TABLE 10.67 -- continued

MODEL #	4	5	6	7
DESCRIPTION	Non-Discriminatory	Discriminatory	Non-Discriminatory	Discriminatory
COMPENSATION ELEMENTS	2	2	6	6
VARIABLES	78	86	94	102
MEAN AGENT FACTORS				
Agents who Quit	221	128	102	-60
Fired Agents	3380	1320	904	570
Normal Agents	38	2620	572	65
All Agents	584	436	185	-20
MEAN SATISFACTION OF AGENTS				
Agents who Quit	-223	-221	-66	-65
Fired Agents	-240	-300	-57	-66
Normal Agents	-189	-166	-50	-64
All Agents	-228	-215	-63	-65
MEAN NUMBER OF INTERACTIONS (NORMALIZED)				
Agents who Quit	2.3597	2.0180	3.9383	3.8724
Fired Agents	5.5558	3.5605	7.5252	5.8540
Normal Agents	2.8700	1.9850	5.7950	5.5150

TABLE 10.67 -- continued

MODEL #	4	5	6	7
DESCRIPTION	Non-Discriminatory	Discriminatory	Non-Discriminatory	Discriminatory
COMPENSATION ELEMENTS	2	2	6	6
VARIABLES	78	86	94	102
All Agents	2.5100	2.0958	4.2363	4.1423
Principal's Satisfaction	-332.2 (129.7)	-406.2 (161.1)	-227.7 (105.5)	-234.6 (121.9)
Principal's Factor	2.8789 (7.6635)	1.8123 (5.5003)	-0.1291 (5.9218)	0.2479 (5.8992)

TABLE 10.68: Probability Distributions for Models 4, 5, 6, and 7

VARIABLE		NOMINAL VALUES (Code Mappings)				
		1	2	3	4	5
AGE, A		≤ 20	(20,25]	(25,35]	(35,55]	> 55
Prob(A)		0.10	0.15	0.30	0.35	0.10
EDUCATION, D		none	high school	vocational	undergrad	graduate
Prob(D A)		1	2	3	4	5
A	1	0.10	0.30	0.40	0.20	0.00
	2	0.10	0.20	0.40	0.20	0.10
	3	0.05	0.10	0.30	0.50	0.05
	4	0.05	0.05	0.30	0.30	0.30
	5	0.00	0.10	0.10	0.30	0.50
EXPERIENCE, X		none	≤ 2 years	≤ 5 years	≤ 20 years	> 20 years
Prob(X A)		1	2	3	4	5
A	1	0.70	0.20	0.10	0.00	0.00
	2	0.60	0.30	0.10	0.00	0.00
	3	0.20	0.40	0.30	0.10	0.00
	4	0.00	0.10	0.30	0.60	0.00
	5	0.00	0.00	0.00	0.20	0.80
GENERAL SOCIAL SKILLS, GSS Prob(GSS A)		1	2	3	4	5
A	1	0.20	0.30	0.30	0.15	0.05
	2	0.10	0.40	0.30	0.10	0.10
	3	0.10	0.20	0.40	0.20	0.10
	4	0.05	0.10	0.20	0.40	0.25
	5	0.05	0.10	0.20	0.30	0.35
OFFICE AND MANAGERIAL SKILLS, OMS Prob(OMS D)		1	2	3	4	5
D	4	0.60	0.25	0.05	0.05	0.05
	2	0.50	0.20	0.15	0.40	0.05
	3	0.30	0.30	0.20	0.10	0.10
	4	0.10	0.10	0.20	0.40	0.20
	5	0.05	0.05	0.30	0.40	0.20

CHAPTER 11

CONCLUSION

The basic model for the study of agency theory includes knowledge bases, behavior and motivation theory, contracts contingent on the behavioral characteristics of the agent, and learning with genetic algorithms.

The initial experiments were aimed at an exploration of the new methodology. The goal was to deal with any technical issues in learning, such as length of simulation, convergence behavior of solutions, choice of values for the genetic operators. The initial studies were motivated by questions of the following nature:

- * Can this new framework be used to tailor contracts to the behavioral characteristics of the agents?
- * Is it worthwhile to include in the contract elements of compensation other than fixed pay and share of output?
- * How can good contracts be characterized and understood?

Model 3 of agency described in detail in Chapter 9, includes and incorporates theories of behavior and motivation, dynamic learning, and complex compensation plans was examined from the viewpoint of different informational assumptions. The results from Model 3 show that the traditional agency models are inadequate for identifying

important elements of compensation plans. The reasons that the traditional models fail are their strong methodological assumptions and lack of a framework which deals with complex behavioral and motivational factors, and their influence in inducing effort selection in the agent. Model 3 attempts to remove this inadequacy.

The results of this research are dependent on the informational assumptions of the principal and the agent. It is not suggested that the traditional theory is always wrong. In some cases (i.e. for the informational assumptions of some principal), both theories may agree on their recommendations for optimal compensation plans. However, this research does present several significant counter-examples to traditional agency wisdom. Sec. 9.5 contains the details.

Models 4 through 7 have comparatively more realism. These models simulate a multi-agent, multi-period, and dynamic agency models which include contracts contingent on the characteristics of the agents. The antecedents are not point estimates as in the earlier studies, but interval estimates. This made it possible to use specialization and generalization operators as learning mechanisms in addition to genetic operators. Further, while these models followed the basic LEN model (as did the previous models), the agents who enter the agency all have different risk aversions and reservation welfares.

Models 4 and 5 have only two elements of compensation each, while Models 6 and 7 have six each. This enables one to study the effect of complex contracts as opposed to simple contracts. Moreover, in Models 5 and 7, the principal evaluates the agents individually. Performance of an agent is not compared to that of the others. As

a consequence, the firing policy is individualized to the agents. This is described as a "discriminatory" policy. In Models 4 and 6, the evaluation of the performance of an agent is relative to the performance of the other agents. Hence, there is one common firing policy for all agents. This policy is described as a "non-discriminatory" policy. This design of the experiments enables one to study the effect of the two policies on agency performance.

Models 4 through 7 reveal several interesting results. The practice of discriminatory evaluation of performance is beneficial to some agents (those who work hard and are well motivated), while it is detrimental to others (shirkers). Discrimination is not a desirable policy for the principal, since the mean satisfaction obtained by the principal in the discriminatory models is comparatively less. However, a discriminatory evaluation may serve to bootstrap an organization having low morale (by firing the shirkers), and ensuring the highest rate of increase of satisfaction for the principal.

Increasing the complexity of contracts ensures low agent turnover (because of increased flexibility) and increased overall satisfaction. This finding takes on added significance when the cost of human resource management (such as hiring, terminating, and training) is taken into account. This is suggested as future research in Chapter 12.

Complexity of contracts and the selective practice of relative evaluation of agent performance are powerful tools which can be used by the principal to achieve the goals of the agency. Their interaction and the trade-offs involved are, however, far from straight-forward. Section 10.4 through 10.8 provide the details. Further research is

necessary to explore these agency mechanisms more fully. Suggestions are given in Chapter 12.

On the one hand, each of the Models 4 through 7 seem to act as templates for organizations with different goals. On the other, a model which accurately reflects an existing organization may be chosen for simulation of the agency environment.

CHAPTER 12 FUTURE RESEARCH

A number of directions for future research are possible. These directions are related to the nature of the agency, behavior and motivation theory, additional learning capabilities, and the role of maximum entropy.

12.1 Nature of the Agency

The following enhancements to the agency attempt to include greater realism. This would enable the study of existing agencies, and would ensure applicability of research results.

1. The principal warns an agent whenever his performance triggers a firing decision. The number of warnings could be a control variable to study agency models.
2. The role of the private information of the agent could be control variable - number of elements of private information, and changing their values in different periods of the agency.
3. Agents modify their behavior with time. This is an extremely realistic situation. This would imply that the agents also employ learning mechanisms. The effort selection mechanism of the agents and their acceptance/rejection criteria would

- be specified by knowledge bases. This makes it possible to apply learning to these knowledge bases, the same as was done for the principal.
4. Inclusion of the cost of human resource management for the principal. This cost might be included either in the computation of rule fitnesses, or in the firing rules of the principal. Coupled with a learning mechanism, this would ensure that the principal learns the correct criteria, and to change the criteria in response to the exogenous environment.
 5. The type of simulation involved in the study of the models is discrete-event. The time from the acceptance of a contract to the sharing of the output is one indestructible time slice. In reality, there is a small probability for an agent to resign, or for the principal to fire an agent, before the completion of the contract period. This may be a desirable extension to the models above, and would be a step towards achieving continuous simulation.

12.2 Behavior and Motivation Theory

As was pointed out in Chapter 9, further research is necessary in order to unveil the cause-effect relationships among the elements of the behavioral models, and their influence in unearthing good contracts in the learning process. This might shed insight into the correlations observed among the various elements of compensation in Model 3. Further research varying the "functional" assumptions must be carried out if any clear pattern is to emerge. This would also help estimate the robustness of the model (the change in the degree and direction of the correlations when the functional specifications

are changed). However, a correlational study of the compensation variables in the final knowledge base is a starting point for characterizing good contracts. The acceptance or rejection of contracts by the agents, or the effort-inducing influence of different contracts, may be better predicted by forming correlational links between the different compensation elements.

One potential benefit in investigating the role of behavior and motivation theory is that compensation rules may be modified according to correlations. For example, if, for a particular class of agents, benefits and share of output are strongly positively correlated, then all rules that do not reflect this property may be discarded. Normal genetic operators may then be applied. The mutation operator would ensure exploration of new rules in the search space, while the correlation-modified rules would fix the rule population in a desirable sector of the search space. This procedure may not be defensible if, upon further research, it was found that the correlations are purely random. This research indicates that this is unlikely to happen.

12.3 Machine Learning

PAC-learning may be applied to the set of final contracts in order to determine their closeness to optimality. Genetic algorithms do not guarantee optimality, even though in practice they perform well. However, some measure of goodness of solutions is necessary. PAC-learning, described in Chapter 2, provides such a measure along with the confidence level. PAC theory is statistical and non-parametric in nature.

The learning mechanisms may also be expanded. For example, as pointed out in Section 12.2 above, learning could be modified by correlational findings if a significant causal relationship could be found between motivation theory and the identification of good contracts.

The genetic operators may be varied in future research. For example, only one-point uniform crossover was used. The number of crossover points could be increased. Similarly, the mutation operator may be made dependent on time (or the number of learning periods). The knowledge base may also be coded as a binary string, instead of being a string of multi-valued nominal and ordinal attributes. Instead of randomly trying all combinations of genetic operators and codings, the structure of the knowledge base should be studied in order to see if there are any clues that point to the superiority of one scheme over the other.

Another interesting, and quite important, research is the study of deceptiveness of the knowledge base. A particular coding of the strings (which are the rules) might yield a population that deceives the genetic algorithm. This implies that the population of strings wanders away from the global optimum. An examination of the learning statistics of Chapter 10 suggest that such deception might be happening in Models 4 through 7. Deceptiveness is characterized as the tendency of hyperplanes in the space of building blocks to direct search in non-optimal directions. The domain of the theory of genetic algorithms is the n -dimensional euclidean space, or its subspaces (such as the n -dimensional binary space). The main problem in the study of deceptiveness in the models used in this research is that the relevant search spaces are n -dimensional nominal

and ordinal valued spaces. It remains to be seen how to adapt the theory of genetic algorithms as dynamical systems to such models. It is encouraging to note that the deterioration in average fitness with increasing learning periods (as in Models 4 through 7) is minor, suggesting that the model might be GA-deceptive instead of GA-hard (GA-hard problems are those that seriously mislead the genetic algorithm). Further encouragement derives from Whitley's theorem which states that the only challenging problems are GA-deceptive (Whitley, L.D., 1991). Hence, one is at least assured, while studying the more realistic models of agency, that these models are in fact sufficiently challenging.

It is fairly straight-forward to include learning mechanisms wherever knowledge bases are employed. In addition to having knowledge bases for selection of appropriate compensation schemes and firing of agents, one may also have knowledge bases for the agent(s) for effort selection, acceptance or rejection of contracts, and resigning from the agency. The rules for calculation of satisfaction or welfare in the agency may be made as extensive and detailed as one pleases. This highlights the flexibility of the new framework - it is possible to extend the model by adding knowledge bases in a modular manner without increasing the complexity of the simulation beyond that caused by the size and number of knowledge bases. In contrast, models in mathematical optimization quickly become intractable by the addition of additional variables.

12.4 Maximum Entropy

Maximum entropy (MaxEnt) distributions seek to capture all the information about a random variable without introducing unwarranted bias in the distribution. This is called "maximal non-committalness". The information of the agents and of the principal is specified by using probability distributions. The role of MaxEnt distributions was not attempted in this thesis. It is worthwhile to pursue the question of whether using a maximum entropy distribution having the same mean and variance as the original distribution makes any difference. In other words, an interesting future study might be examination of the "MaxEnt robustness" of agency models. The results might have interesting implications. If the results show that agency models coupled with learning (as in this thesis) are MaxEnt robust, then it is not necessary to insist on using MaxEnt distributions (which are computationally difficult to find). Similarly, if the models are not MaxEnt robust, then deviation from MaxEnt behavior might yield a clue about the tradeoffs involved in seeking a MaxEnt distribution.

APPENDIX FACTOR ANALYSIS

We use the SAS procedures (PROC FACTOR) which uses Principal Components Method to extract factors from the final rule population (Guertin and Bailey, 1970). We also subject the data to Kaiser's Varimax rotation, which is employed to avoid skewed distribution of variance explained by the factors. In the initial "direct" solution, the first factor accounts for most of the variance, followed in decreasing order by the rest of the factors. In the "derived" solution (i.e. after rotation), variables load either maximum or close to zero. This enables the factors to stand out more sharply. By the Kaiser criterion, factors whose eigenvalues are greater than one are retained since they are deemed to be significant. The size of a sample (or the size of the population) approximates the degrees of freedom for testing significance of factor loadings. Using 500 degrees of freedom (the population size) and a relatively stringent 1% significance level, the critical value of correlation is 0.115. The critical value of a factor loading, f_c , is given by the Burt-Banks formula (Child, 1990):

$$f_c = r_c \sqrt{\frac{n}{n+1-r}},$$

where r_c is the critical correlation value, n is the number of variables, and r is the position number of the factor being considered.

The Burt-Banks formula ensures that the acceptable level of factor loadings increases for later factors, so that the criteria for significance become more stringent as one progresses from the first factor to higher factors. This is essential, because specific variance plays an increasing role in later factors at the expense of common variance. The Burt-Banks formula, in addition to adjusting the significance, also accounts for the sample size and the number of variables.

REFERENCES

- Alchian, A.A. and Demsetz, H. (1972). "Production, Information Costs and Economic Organization." *American Economic Review* 62(5), pp. 777-795.
- Anderson, J.R. (1976). *Language, Memory, and Thought*. Lawrence Erlbaum, Hillsdale, N.J.
- Anderson, J.R. (1980). *Cognitive Psychology and its Implications*. W.H. Freeman and Co., San Francisco.
- Anderson, J.R., and Bower, G.H. (1973). *Human Associative Memory*. Winston, Washington, D.C.
- Angluin, D. (1987). "Learning k-term DNF Formulas Using Queries and Counterexamples." Technical Report YALEU/DCS/RR-559, Yale University.
- Angluin, D. (1988). "Queries and Concept Learning." *Machine Learning* 2, pp. 319-342.
- Angluin, D., and Laird, P. (1988). "Learning from Noisy Examples." *Machine Learning* 2, pp. 343-370.
- Angluin, D., and Smith, C. (1983). "Inductive Inference: Theory and Methods." *ACM Comp. Surveys* 15(3), pp. 237-270.
- Arrow, K.J. (1986). "Agency and the Market." In *Handbook of Mathematical Economics III*, Chapter 23; Arrow, K.J., and Intriligator, M.D. (eds.), North-Holland, Amsterdam, pp. 1183-1200.
- Bamberg, G., and Spremann, K. (Eds.) (1987). *Agency Theory, Information, and Incentives*. Springer-Verlag, Berlin.
- Baron, D. (1989). "Design of Regulatory Mechanisms and Institutions." In *Handbook of Industrial Organization II*, Chap. 24; Schmalensee, R. and Willig, R.D. (eds.), Elsevier Science Publishers, New York.

- Barr, A., Cohen, P.R., and Feigenbaum, E.A. (1989). *The Handbook of Artificial Intelligence IV*. Addison-Wesley Publishing Company, Reading, MA.
- Barr, A., and Feigenbaum, E.A. (1981). *The Handbook of Artificial Intelligence I*. William Kaufman, Los Altos, CA.
- Barr, A., and Feigenbaum, E.A. (1982). *The Handbook of Artificial Intelligence II*. William Kaufman, Los Altos, CA.
- Berg, S., and Tschirhart, J. (1988a). *Natural Monopoly Regulation: Principles and Practice*. Cambridge University Press, New York.
- Berg, S., and Tschirhart, J. (1988b). "Factors Affecting the Desirability of Traditional Regulation." Working Paper, Public Utilities Research Center, University of Florida, Gainesville, FL.
- Besanko, D., and Sappington, D. (1987). *Designing Regulatory Policy with Limited Information*. Harwood Academic Publishers, London.
- Blickle, M. (1987). "Information Systems and the Design of Optimal Contracts." In *Agency Theory, Information, and Incentives*; Bamberg, G. and Spremann, K. (eds.), Springer-Verlag, Berlin, pp. 93-103.
- Blumer, A., Ehrenfeucht, A., Haussler, D., and Warmuth, M.K. (1989). "Learnability and the Vapnik-Chervonenkis Dimension." *JACM* 36(4), pp. 929-965.
- Brown, S.J., and Sibley, D.S. (1986). *The Theory of Public Utility Pricing*. Cambridge University Press, New York.
- Buchanan, B.G., and Feigenbaum, E.A. (1978). "DENDRAL and META-DENDRAL: Their Application Dimension." *Artificial Intelligence* 11, pp. 5-24.
- Buchanan, B.G., Mitchell, T.M., Smith, R.G., and Johnson, C.R. Jr. (1977). "Models of learning systems." In Belzer, J., Holzman, A.G., and Kent, A. (eds.), *Encyclopedia of Computer Science and Technology* 11, Marcel Dekker, New York, pp. 24-51.
- Campbell, J.P., and Pritchard, R.D. (1976). "Motivation Theory in Industrial and Organizational Psychology." In *Handbook of Industrial and Organizational Psychology*; Dunnette, M. (ed.), Rand McNally, Chicago.
- Cannon, W.B. (1939). *The Wisdom of the Body*. Norton, New York.
- Child, D. (1990). *The Essentials of Factor Analysis*. Cassell, London.

- Christensen, J. (1981). "Communication in Agencies." *The Bell Journal of Economics* 12, pp. 661-674.
- Cohen, P.R., and Feigenbaum, E.A. (1982). *The Handbook of Artificial Intelligence III*. William Kaufman, Los Altos, CA.
- Conway, W. (1986). "Application of Decision Analysis to New Product Development - A Case Study." In *Computer Assisted Decision Making*, Mitra, G. (ed.), North-Holland, New York.
- Craik, F.I.M., and Tulving, E. (1975). "Depth of Processing and the Retention of Words in Episodic Memory." *Journal of Experimental Psychology: General* 104, pp. 268-294.
- De Jong, K.A. (1988). "Learning with Genetic Algorithms: an Overview." *Machine Learning* 3(2), pp. 121-138.
- Demsetz, H. (1968). "Why Regulate Utilities?" *Journal of Law and Economics* 7, pp. 55-65.
- Demski, J., and Sappington, D.E.M. (1984). "Optimal Incentive Schemes with Multiple Agents." *Journal of Economic Theory* 33, pp. 152-171.
- Dietterich, T.G., and Michalski, R.S. (1979). "Learning and Generalization of Characteristic Descriptions: Evaluation Criteria and Comparative Review of Selected Methods." *Artificial Intelligence* 16, pp. 257-294.
- Dietterich, T.G., and Michalski, R.S. (1983). "A Comparative Review of Selected Methods from Learning from Examples." In *Machine Learning: An Artificial Intelligence Approach*, Morgan Kaufmann, San Mateo, CA.
- Duda, R.O., Hart, P.E., Konolige, K., and Reboh, R. (1979). "A Computer-based Consultant for Mineral Exploration." Technical Report, SRI International.
- Dudley, R.M. (1978). "Central Limit Theorems for Empirical Measures." *Annals of Probability* 6(6), pp. 899-929.
- Dudley, R.M. (1984). "A Course on Empirical Processes." *Lecture Notes in Mathematics* 1097, pp. 2-142.
- Dudley, R.M. (1987). "Universal Donsker Classes and Metric Entropy." *Annals of Probability* 15(4), pp. 1306-1326.

- Einhorn, H.J. (1982). "Learning from Experience and Suboptimal Rules in Decision Making." In *Judgment Under Uncertainty: Heuristics and Biases*; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 268-286.
- Ellig, B.R. (1982). *Executive Compensation - A Total Pay Perspective*. McGraw-Hill, New York.
- Engelberger, T.F. (1980). *Robotics in Practice*. Kegan Paul, London.
- Erman, L.D., Hayes-Roth, F., Lesser, V.R., and Reddy, D.R. (1980). "The Hearsay-II Speech Understanding System: Integrating Knowledge to Resolve Uncertainty." *Computing Surveys* 12(2).
- Erman, L.D., London, P.E., and Fickas, S.F. (1981). "The Design and an Example Use of Hearsay III." *Proc. IJCAI* 7.
- Feigenbaum, E.A., and McCorduck, P. (1983). *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*. Addison-Wesley, Reading, MA.
- Firchau, V. (1987). "Information Systems for Principal-Agent Problems." In *Agency Theory, Information, and Incentives*; Bamberg, G. and Spremann, K. (eds.), Springer-Verlag, Berlin, pp. 81-92.
- Fishburn, P.C. (1981). "Subjective Expected Utility: A Review of Normative Theories." *Theory and Decision* 13, pp. 139-199.
- Gjesdal, F. (1982). "Information and Incentives: The Agency Information Problem." *Review of Economic Studies* 49, pp. 373-390.
- Glass, A.L., and Holyoak, K.J. (1986). *Cognition*. Random House, New York, NY.
- Grandy, C. (1991). "The Principle of Maximum Entropy and the Difference Between Risk and Uncertainty." In *Maximum Entropy and Bayesian Methods*; Grandy, W.T. Jr. and Schick, L.H. (eds.), Kluwer Academic Publishers, Boston, pp. 39-47.
- Grossman, S.J. and Hart, O.D. (1983). "An Analysis of the Principal-Agent Problem." *Econometrica* 51(1), pp. 7-45.
- Guertin, W.H., and Bailey, J.P. Jr. (1970). *Introduction to Factor Analysis*. Guertin, W.H., and Bailey, J.P. Jr.

- Harris, M. and Raviv, A. (1979). "Optimal Incentive Contracts with Imperfect Information." *Journal of Economic Theory* 20, pp. 231-259
- Hart, P.E., Duda, R.O., and Einaudi, M.T. (1978). "A Computer-based Consultation System for Mineral Exploration." Technical Report, SRI International.
- Hasher, L., and Zacks, R.T. (1979). "Automatic and Effortful Processes in Memory." *Journal of Experimental Psychology: General* 108, pp. 356-388.
- Haussler, D. (1988). "Quantifying Inductive Bias: AI Learning Algorithms and Valiant's Learning Framework." *Artificial Intelligence* 36, pp. 177-221.
- Haussler, D. (1989). "Learning Conjunctive Concepts in Structural Domains." *Machine Learning* 4, pp. 7-40.
- Haussler, D. (1990a). "Applying Valiant's Learning Framework to AI Concept-Learning Problems." In *Machine Learning: An Artificial Intelligence Approach*, Vol. III; Kodratoff, Y. and Michalski, R. (eds.), Morgan Kaufmann, San Mateo, CA., pp. 641-669.
- Haussler, D. (1990b). "Decision Theoretic Generalizations of the PAC Learning Model for Neural Net and Other Learning Applications." Technical Report UCSC-CRL-91-02, University of CA, Santa Cruz.
- Hayes-Roth, F., and Lesser, V.R. (1977). "Focus of Attention in the Hearsay-II System." *Proc. IJCAI* 5.
- Hayes-Roth, F., and McDermott, J. (1978). "An Interference Matching Technique for Inducing Abstractions." *CACM* 21(5), pp. 401-410.
- Hebb, D.O. (1961). "Distinctive Features of Learning in the Higher Animal." In *Brain Mechanisms and Learning*; Delafresnaye, J.F. (ed.), Blackwell, London.
- Holland, J.H. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI.
- Holland, J.H. (1986). "Escaping Brittleness: The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-Based Systems." In *Machine Learning: An Artificial Intelligence Approach 2*; Michalski, R.S., Carbonell, J.G., and Mitchell, T.M. (eds.), Morgan Kaufmann, Los Altos, CA, pp. 593-623.
- Holmstrom, B. (1979). "Moral Hazard and Observability." *Bell Journal of Economics* 10, pp. 74-91.

- Holmstrom, B. (1982). "Moral Hazard in Teams." *Bell Journal of Economics* 13, pp. 324-340.
- Hull, C.L. (1943). *Principles of Behavior*. Appleton-Century-Crofts, New York.
- Jaynes, E.T. (1982). "On the Rationale of Maximum-Entropy Methods." *Proc. IEEE* 70(9).
- Jaynes, E.T. (1983). *Papers on Probability, Statistics, and Statistical Physics: A Reprint Collection*. Rosenkrantz, R.D. (ed.). North-Holland, Amsterdam.
- Jaynes, E.T. (1986a). "Bayesian Methods: General Background -An Introductory Tutorial." In *Maximum Entropy and Bayesian Methods in Applied Statistics*; Justice, J.H. (ed.), Cambridge University Press, New York, pp. 1-25.
- Jaynes, E.T. (1986b). "Monkeys, Kangaroos, and N." In *Maximum Entropy and Bayesian Methods in Applied Statistics*; Justice, J.H. (ed.), Cambridge University Press, New York, pp. 26-58.
- Jaynes, E.T. (1991). "Notes on Present Status and Future Prospects." In *Maximum Entropy and Bayesian Methods*, Grandy, W.T. Jr. and Schick, L.H. (eds.), Kluwer Academic Publishers, Boston, MA, pp. 1-13.
- Jensen, M.C., and Meckling, W.H. (1976). "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure." *Journal of Financial Economics* 3, pp. 305-360.
- Kahn, A.E. (1978). "Applying Economics to an Imperfect World." *Regulation*, pp. 17-27.
- Kahneman, D., and Tversky, A. (1982a). "Subjective Probability: A Judgment of Representativeness." In *Judgment Under Uncertainty: Heuristics and Biases*; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 32-47.
- Kahneman, D., and Tversky, A. (1982b). "On the Psychology of Prediction." In *Judgment Under Uncertainty: Heuristics and Biases*; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 48-68.
- Kahneman, D., and Tversky, A. (1982c). "On the Study of Statistical Intuitions." In *Judgment Under Uncertainty: Heuristics and Biases*; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 493-508.

- Kahneman, D., and Tversky, A. (1982d). "Variants of Uncertainty." In *Judgment Under Uncertainty: Heuristics and Biases*; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 509-520.
- Kane, T.B. (1991). "Reasoning with Maximum Entropy in Expert Systems." In *Maximum Entropy and Bayesian Methods*; Grandy, W.T. Jr. and Schick, L.H. (eds.), Kluwer Academic Publishers, Boston, MA, pp. 201-213.
- Keeney, R.L. (1984). *Decision Analysis: An Overview*. Wiley, New York.
- Keeney, R.L., and Raiffa, H. (1976). *Decisions with Multiple Objectives*. Wiley, New York.
- Kodratoff, Y., and Michalski, R. (1990). *Machine Learning: An Artificial Intelligence Approach III*. Morgan Kaufmann, San Mateo, CA.
- Kolmogorov, A.N., and Tihomirov, V.M. (1961). " ϵ -Entropy and ϵ -Capacity of Sets in Functional Spaces." *American Mathematical Society Translations (Series 2)* 17, pp. 277-364.
- Kullback, S. (1959). *Information Theory and Statistics*. Wiley, New York.
- Lenat, D.B. (1977). "On Automated Scientific Theory Formation: A Case Study Using the AM Program." In *Machine Intelligence 9*; Hayes, J.E., Michie, D.M., and Mikulich, L.I. (eds.), Halstead Press, New York, pp. 251-286.
- Lewis, T.R., and Sappington, D.E.M. (1991a). "Should Principals Inform Agents?: Information Management in Agency Problems." Working Paper, Department of Economics, University of Florida Mimeo, Gainesville, FL.
- Lewis, T.R., and Sappington, D.E.M. (1991b). "Selecting an Agent's Ability." Working Paper, Department of Economics, University of Florida Mimeo, Gainesville, FL.
- Lichtenstein, S., Baruch, F., and Phillips, L.D. (1982). "Calibration of Probabilities: The State of the Art to 1980." In *Judgment Under Uncertainty: Heuristics and Biases*; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 306-334.
- Lindley, D.V. (1971). *Making Decisions*. Wiley, New York.
- Lippman, A. (1988). "A Maximum Entropy Method for Expert System Construction." In *Maximum Entropy and Bayesian Methods in Science and Engineering 2*;

- Erickson, G.J. and Smith, C.R. (eds.), Kluwer Academic Publishers, Boston, MA, pp. 243-263.
- Mandelbrot, B.B. (1982). *The Fractal Geometry of Nature*. W. H. Freeman, San Francisco.
- Mandler, G. (1967). "Organization and Memory." In *The Psychology of Learning and Motivation I*; Spence, K.W., and Spence, J.T., (eds.), Academic Press, New York, pp. 327-372.
- Maslow, A.H. (1943). "A Theory of Human Motivation." *Psychological Review* 50, pp. 370-396.
- Maslow, A.H. (1954). *Motivation and personality*. Harper, New York.
- McDermott, D., and Forgy, C. (1978). "Production System Conflict Resolution Strategies." In *Pattern Directed Inference Systems*; Waterman, D.A. and Hayes-Roth, F. (eds.), Academic Press, New York, pp. 177-199.
- Michalski, R. (1983). "A Theory and Methodology of Inductive Learning." In *Machine Learning: An Artificial Intelligence Approach*; Michalski, R., Carbonell, J.G., and Mitchell, T.M. (eds.), Morgan Kaufmann, San Mateo, CA, pp. 83-134.
- Michalski, R., Carbonell, J.G., and Mitchell, T.M. (1983). *Machine Learning*. Tioga, Palo Alto, CA.
- Minsky, M. (1975). "A Framework for Representing Knowledge." In *The Psychology of Computer Vision*; Winston, P. (ed.), McGraw-Hill, New York.
- Mitchell, T.R. (1974). "Expectancy Models of Job Satisfaction, Occupational Preference and Effort: A Theoretical, Methodological, and Empirical Approach." *Psychological Bulletin* 81, pp. 1053-1077.
- Mitchell, T.M. (1977). "Version Spaces: A Candidate Elimination Approach to Rule Learning." *IJCAI* 5, pp. 305-310.
- Mitchell, T.M. (1979). "An Analysis of Generalization as a Search Problem." *IJCAI* 6, pp. 577-582.
- Mitchell, T.M. (1982). "Generalization as Search." *Artificial Intelligence* 18, pp. 203-226.
- Mitra, G. (ed.) (1986). *Computer Assisted Decision Making*. North-Holland, New York.

- Newell, A., and Simon, H.A. (1956). "The Logic Theory Machine." IRE Transactions on Information Theory 2, pp. 61-79.
- Nilsson, N.J. (1980). Principles of Artificial Intelligence. Tioga, Palo Alto, CA.
- Nisbett, R.E., Borgida, E., Crandall, R., and Reed, H. (1982). "Popular Induction: Information is Not Necessarily Informative." In Judgment Under Uncertainty: Heuristics and Biases; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 101-116.
- Norman, D.A. and Rumelhart, D.E. (1975). Explorations in Cognition. Freeman, San Francisco.
- Paul, R.P. (1981). Robot Manipulators: Mathematics, Programming and Control. MIT Press, Cambridge, MA.
- Peikoff, L. (1991). Objectivism: The Philosophy of Ayn Rand. Dutton Books, New York.
- Phillips, L.D. (1986). "Decision Analysis and its Application in Industry." In Computer Assisted Decision Making; Mitra, G. (ed.), North-Holland, New York.
- Pitt, L., and Valiant, L.G. (1988). "Computational Limitations on Learning from Examples." JACM 35(4), pp. 965-984.
- Pollard, D. (1984). Convergence of Stochastic Processes. Springer-Verlag, New York.
- Porter, L.W., and Lawler, E.E. (1968). Managerial Attitudes and Performance. Irwin-Dorsey, Homewood, IL.
- Pratt, J.W. and Zeckhauser, R.J. (1985). Principals and Agents: The Structure of Business. Harvard University Press, Boston, MA.
- Quillian, M.R. (1968). "Semantic Memory." In Semantic Information Processing; Minsky, M. (ed.), MIT Press, Cambridge, MA, pp. 216-270.
- Quinlan, J.R. (1979). "Discovering Rules from Large Collections of Examples: A Case Study." In Expert Systems in Microelectronic Age; Michie, D. (ed.), Edinburgh University Press, Edinburgh.
- Quinlan, J.R. (1986). "Induction of Decision Trees." Machine Learning 1(1), pp. 81-106.

- Quinlan, J.R. (1990). "Probabilistic Decision Trees." In *Machine Learning: An Artificial Intelligence Approach III*; Kodratoff, Y. and Michalski, R. (eds.), Morgan Kaufmann, San Mateo, CA.
- Rand, A. (1967). *Introduction to Objectivist Epistemology*. Mentor Books, New York.
- Rasmusen, E. (1989). *Games and Information: An Introduction to Game Theory*. Basil Blackwell, New York.
- Reber, A.S. (1967). "Implicit Learning of Artificial Grammars." *Journal of Verbal Learning and Verbal Behavior* 5, pp. 855-863.
- Reber, A.S. (1976). "Implicit Learning of Synthetic Languages: The Role of Instructional Set." *Journal of Experimental Psychology: Human Learning and Memory* 2, pp. 88-94.
- Reber, A.S., and Allen, R. (1978). "Analogy and Abstraction Strategies in Synthetic Grammar Learning: A Functional Interpretation." *Cognition* 6, pp. 189-221.
- Reber, A.S., Kassin, S.M., Lewis, S., and Cantor, G.W. (1980). "On the Relationship Between Implicit and Explicit Modes in the Learning of a Complex Rule Structure." *Journal of Experimental Psychology: Human Learning and Memory* 6, pp. 492-502.
- Rees, R. (1985). "The Theory of Principal and Agent - Part I." *Bulletin of Economic Research* 37(1), pp. 3-26.
- Riesbeck, C.K., and Schank, R.C. (1989). *Inside Case-Based Reasoning*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Rivest, R. (1987). "Learning Decision-Lists." *Machine Learning* 2(3), pp. 229-246.
- Ross, S.A. (1973). "The Economic Theory of Agency: The Principal's Problem." *American Economic Review* 63(2), pp. 134-139.
- Samuel, A.L. (1963). "Some Studies in Machine Learning Using the Game of Checkers." In *Computers and Thought*; Feigenbaum, E. A. and Feldman, J. (eds.), McGraw-Hill, New York, pp. 71-105.
- Sammut, C., and Banerji, R. (1986). "Learning Concepts by Asking Questions." In *Machine Learning: An Artificial Intelligence Approach II*; Michalski, R., Carbonell, J.G., and Mitchell, T.M. (eds.), Morgan Kaufmann, San Mateo, CA.

- Sappington, D.E.M. and Sibley, D.S. (1988). "Regulating Without Cost Information: the Incremental Surplus Subsidy Scheme." *International Economic Review* 29(2).
- Sappington, D.E.M. and Stiglitz, J.E. (1987). "Information and Regulation." In *Public Regulation*; Bailey, E.E. (ed.), MIT Press, Cambridge, MA.
- Savage, L. (1954). *The Foundations of Statistics*. Wiley, New York.
- Schank, R.C. (1972). "Conceptual Dependency: A Theory of Natural Language Understanding." *Cognitive Psychology* 3, pp. 552-631.
- Schank, R.C., and Abelson, R.P. (1977). *Scripts, Plans, Goals, and Understanding*. Lawrence Erlbaum, Hillsdale, N.J.
- Schank, R.C., and Colby, K.M. (eds.). (1973). *Computer Models of Thought and Language*. Freeman, San Francisco, CA.
- Schneider, D. (1987). "Agency costs and Transaction Costs: Flops in the Principal-Agent Theory of Financial Markets." In *Agency Theory, Information, and Incentives*; Bamberg, G., and Spremann, K.(eds.), Springer-Verlag, Berlin.
- Shavell, S. (1979a). "Risk-sharing and Incentives in the Principal-Agent Relationship." *Bell Journal of Economics* 10, pp. 55-73.
- Shavell, S. (1979b). "On Moral Hazard and Insurance." *Quarterly Journal of Economics* 93, pp. 541-562.
- Shortliffe, E.H. (1976). *Computer-based Medical Consultation: MYCIN*. Elsevier, New York.
- Simon, H.A. (1951). "A Formal Theory of the Employment Relationship." *Econometrica* 19, pp. 293-305.
- Singh, N. (1985). "Monitoring and Hierarchies: The Marginal Value of Information in a Principal Agent Model." *Journal of Political Economy* 93(3), pp. 599-609.
- Spremann, K. (1987). "Agent and Principal." In *Agency Theory, Information, and Incentives*; Bamberg, G., and Spremann, K. (eds.), Springer-Verlag, Berlin, pp. 3-37.
- Steers, R.M., and Porter, L.W. (1983). *Motivation and Work Behavior*. McGraw-Hill, New York.

- Stefik, M., Aikins, J., Balzer, R., Benoit, J., Birnbaum, L., Hayes-Roth, F., and Sacerdoti, E.D. (1982). "The Organization of Expert Systems." *Artificial Intelligence* 18, pp. 135-173.
- Stevens, A.L., and Collins, A. (1977). "The Goal Structure of a Socratic Tutor." BBN Rep. No. 3518, Bolt Beranek and Newman, Inc., Cambridge, MA.
- Stiglitz, J.E. (1974). "Risk Sharing and Incentives in Sharecropping." *Review of Economic Studies* 41, pp. 219-256.
- Tversky, A., and Kahneman, D. (1982a). "Judgment Under Uncertainty: Heuristics and Biases." In *Judgment Under Uncertainty: Heuristics and Biases*; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 3-20.
- Tversky, A., and Kahneman, D. (1982b). "Belief in the Law of Small Numbers." In *Judgment Under Uncertainty: Heuristics and Biases*; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 23-31.
- Tversky, A., and Kahneman, D. (1982c). "Availability: A Heuristic for Judging Frequency and Probability." In *Judgment Under Uncertainty: Heuristics and Biases*; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 163-178.
- Tversky, A., and Kahneman, D. (1982d). "The Simulation Heuristic." In *Judgment Under Uncertainty: Heuristics and Biases*; Kahneman, D., Slovic, P., and Tversky, A., (eds.), Cambridge University Press, New York, pp. 201-210.
- Valiant, L.G. (1984). "A Theory of the Learnable." *CACM* 27 (11), pp. 1134-1142.
- Valiant, L.G. (1985). "Learning Disjunctions of Conjunctions." *Proc. 9th IJCAI* 1, pp. 560-566.
- Vapnik, V.N. (1982). *Estimation of Dependences Based on Empirical Data*. Springer-Verlag, New York.
- Vapnik, V.N., and Chervonenkis, A.Ya. (1971). "On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities." *Theory of Probability and its Applications* 16(2), pp. 264-280.
- Vere, S.A. (1975). "Induction of Concepts in the Predicate Calculus." *Proc. 4th IJCAI*, pp. 281-287.

- Vogelsang, I., and Finsinger, J. (1979). "A Regulatory Adjustment Process for Optimal Pricing by Multiproduct Monopoly Firms." *Bell Journal of Economics* 10, pp. 157-171.
- Weiss, S., and Kulikowski, C. (1991). *Computer Systems that Learn*. Morgan Kaufmann, San Mateo, CA.
- Williams, A.J. (1986). "Decision Analysis for Computer Strategy Planning." In *Computer Assisted Decision Making*, Mitra, G. (ed.), North-Holland, New York.
- Winograd, T. (1972). *Understanding Natural Language*. Academic Press, New York.
- Winograd, T. (1973). "A Procedural Model of Language Understanding." In *Computer Models of Thought and Language*; Schank, R.C., and Colby, K.M. (eds.), Freeman, San Francisco, CA.
- Winston, P. (1975). "Learning Structural Descriptions from Examples." In *The Psychology of Computer Vision*; Winston, P. (ed.), McGraw-Hill, New York.
- Whitley, L.D. (1991). "Fundamental Principles of Deception in Genetic Search." In *Foundations of Genetic Algorithms*; Rawlins, G.J.E. (ed.), Morgan Kaufmann, San Mateo, CA, pp. 221-241.
- Zellner, A. (1991). "Bayesian Methods and Entropy in Economics and Econometrics." In *Maximum Entropy and Bayesian Methods*; Grandy, W.T. Jr., and Schick, L.H. (eds.), Kluwer Academic Publishers, Boston, MA, pp. 17-31.

BIOGRAPHICAL SKETCH

Kiran K. Garimella holds a Master of Computer Applications (M.C.A.) degree from the University of Hyderabad, India (1983-1986) and a Bachelor of Science (with Honors) degree from New Science College, Osmania University, Hyderabad, India (1980-1983). His undergraduate major was chemistry with specialization in biochemistry. His M.C.A. concentration was artificial intelligence and machine learning.


He worked as a software engineer for two years (1986-1988) at Frontier Information Technologies Pvt. Ltd. in Hyderabad, India. His work involved design and development of application software and systems analysis and design studies. He has also consulted with several small businesses in Hyderabad, helping them with computerizing their operations and in the selection of appropriate hardware and applications software. During this time, he was also a part-time doctoral student at the University of Hyderabad, engaged in machine learning research in the Department of Mathematics and Computer Science. He was a guest lecturer at the Institute of Hotel Management, Catering Technology, and Applied Nutrition of the Advanced Training Institute and at the Indian Institute of Computer Science, both in Hyderabad, India. He taught discrete-event system simulation (QMB 4703 - Managerial Operations Analysis III) at the University of Florida, Gainesville, in the Summer A terms of 1990 and 1993

and has been a teaching assistant for information systems, operations research and statistics at the University of Florida.

He secured distinction and first place in the undergraduate class (1982-1983), a University Merit Fellowship (1984), distinction in graduate studies (1985-1986), and the Junior Doctoral Fellowship of the University Grants Commission, India (1987). He holds honorary membership in the Alpha Chapter of Beta Gamma Sigma (1993), and in Alpha Iota Delta (1993).


He has three conference publications (including one book reprint) and is a member of the Association for Computing Machinery, the Decision Sciences Institute, and the Institute of Management Sciences.

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.




Gary J. Koehler, Chairman
Professor of Decision and
Information Sciences

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.




David E. M. Sappington
Lanzilotti-McKethan Professor of
Economics

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.



Richard A. Elnicki
Professor of Decision and
Information Sciences

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.



Anthal Majthay
Associate Professor of
Decision and Information
Sciences

This dissertation was submitted to the Graduate Faculty of the Department of Decision and Information Sciences in the College of Business Administration and to the Graduate School and was accepted as partial fulfillment of the requirements for the degree of Doctor of Philosophy.

August 1993

Dean, Graduate School